Scalable Control Design for Networked Systems

Coordination Through Local Cooperation

Hansson, Jonas

2025

Scalable Control Design for Networked Systems

Coordination Through Local Cooperation

JONAS HANSSON
DEPARTMENT OF AUTOMATIC CONTROL | LUND UNIVERSITY

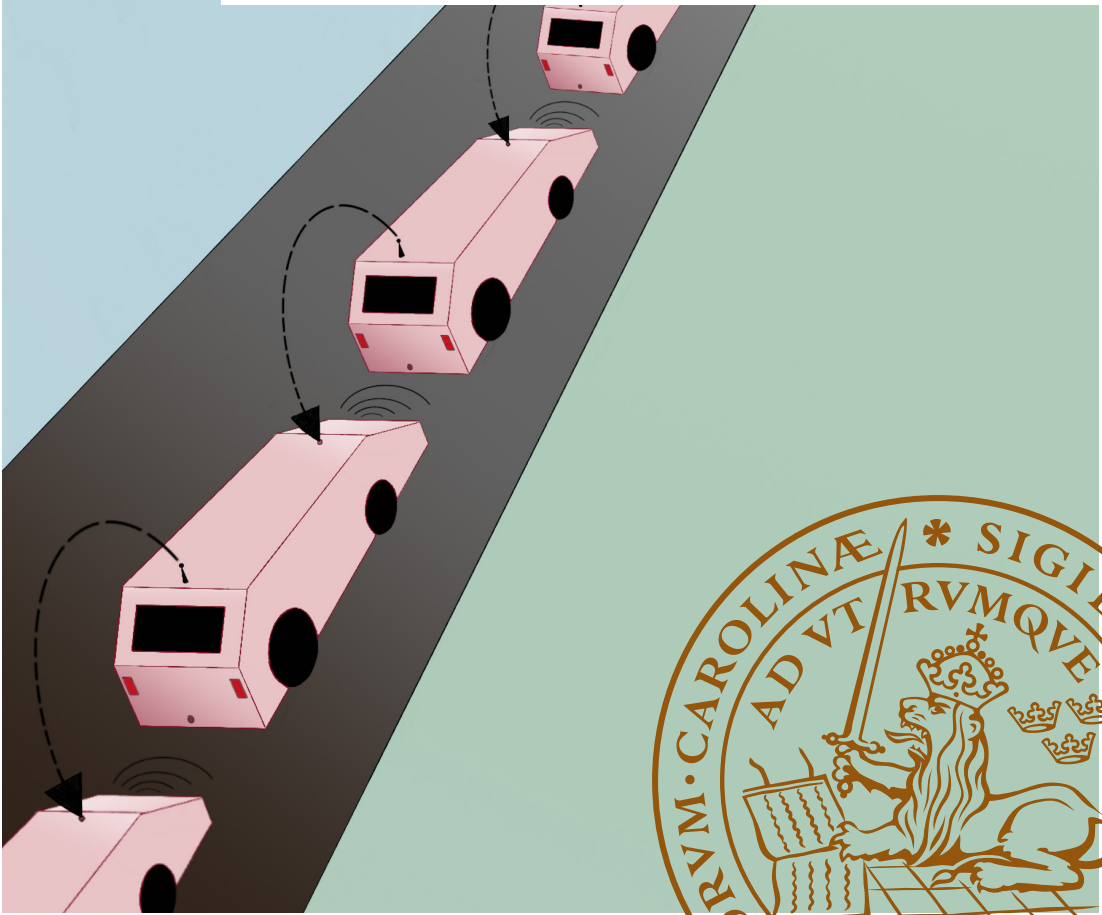

Scalable Control Design for Networked Systems

Coordination Through Local Cooperation

Jonas Hansson

LUND
UNIVERSITY

Department of Automatic Control

The cover page illustrations was made by Linnéa Jönsson.
Printed with permission.

PhD Thesis TFRT-1149
ISBN 978-91-8104-543-7 (print)
ISBN 978-91-8104-544-4 (web)
ISSN 0280–5316

Department of Automatic Control
Lund University
Box 118
SE-221 00 LUND
Sweden

Printed in Sweden by Media-Tryck.
Lund 2025

Abstract

This thesis investigates scalable control design for networked dynamical systems, which are of great importance due to their wide range of practical applications, including large-scale formation control. A central challenge in such systems is enabling agents to coordinate effectively based only on local and relative information, particularly as the system size increases. To address this, the thesis develops a framework for designing robust and scalable coordination protocols that maintain performance even in large formations.

Paper I introduces serial consensus, a novel coordination protocol where the closed-loop system mimics the behavior of multiple simpler consensus protocols interconnected in series. This structure enhances robustness and stability. Paper II applies serial consensus to vehicular formations, demonstrating that it ensures string stability regardless of underlying communication topology. In Paper III, the scalable performance results are extended to the case of high-order coordination of agents with n^{th}-order dynamics ($n \geq 3$). A consequence of this result is that formations of double integrators with local integral control can maintain scalable performance while rejecting constant load disturbances.

Paper IV proposes a class of scalable nonlinear consensus protocols based on the core idea of serial consensus, that is, serializing coordination. This enables coordination of a broader class of systems more closely linked to real-world applications. A unifying theme across Papers I–IV is the study of large-scale dynamical systems and their transient behavior—often in cases where traditional analysis would deem them unstable. To address this, Paper V develops new tools for performance analysis based on ε-pseudospectra, providing a theoretical foundation for understanding transient amplification and robustness in high-order networked systems. Together, these results contribute to a new understanding of design principles of scalable control and its implications for large-scale dynamical systems, offering new perspectives on stability, performance, and robustness in distributed control design.

Acknowledgements

First and foremost, I want to thank Emma—my advisor and academic mentor. From spotting open problems to sticking with me through all the chaotic late-night polishing sessions, your support has been nothing short of incredible. You've always been there with new ideas, a keen eye for structure and detail, and unwavering encouragement—especially when I've had doubts. I couldn't have asked for a better guide through this doctoral journey.

I'm also grateful to the faculty of the Control Department in Lund for fostering such a supportive research environment. In particular, thank you Anders Rantzer, Bo Bernhardsson, and Richard Pates for always being open to questions—whether during a fika chat or a longer blackboard discussion, your input helped me move forward whenever I was stuck. I also want to thank Bassam Bamieh and Francesco Bullo for hosting me during my stay in Santa Barbara. That period had a significant impact on my thinking and shaped many of the directions I've since pursued.

To the admin and technical staff: you are the heroes of the department. Need travel support? Sorted. Computer crashed? Handled. Just need someone to share a fika and a laugh with? Always there. You make it possible for all of us to focus on our research without having to stress about the rest and for this, I'm truly thankful.

To my fellow PhD students and postdocs—thank you for making these years so much fun. From floorball and movie nights to Mello parties, board games, crosswords, running sessions, and all the other spontaneous events—you've been the best part of the day more times than I can count. A special shout-out to Felix, David, and Olle, with whom I've had the great pleasure of sharing many whiteboard sessions.

There are also places that deserve a mention, simply because they've become part of the story: Cancún, Singapore, Paris, Santa Barbara, Norrköping, Louvain-la-Neuve, Hamburgö, the Alps, Nytorpet, Amsterdam, and Cambridge. Whether it was skiing in Austria, adventures in Asia, a karaoke night in California, or road-tripping in Mexico—thank you to everyone who helped make those memories.

Finally, to my family and to Linnéa—thank you for your endless support, love, and patience. Linnéa, you've endured my late-night stress, my pre-deadline crunches, and my post-submission collapses of relief. Thank you for standing by me through all the ups and downs, and for helping me stay sane. I couldn't have done this without you. All love.

Acknowledgement of Funding

The author is with the Department of Automatic Control and the ELLIIT Strategic Research Area at Lund University, Lund, Sweden. This work was partially funded by Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

Contents

1

Introduction

Consensus refers to the process by which multiple agents in a networked system reach agreement on a shared state using only local information. This phenomenon arises in a wide range of systems—from natural flocking behaviors in birds and fish to engineered systems such as vehicle formations, power grid synchronization, and drone swarms. A central challenge in such settings is that agents typically have access only to local measurements. This raises a fundamental question: under what conditions can local interactions guarantee stable, efficient, and robust global coordination, particularly as the system size increases?

To illustrate this, consider a simple vehicle formation in which each vehicle obeys Newton's second law:

$$m_i \ddot{x}_i(t) = u_i(t,x), \tag{1.1}$$

where m_i denotes the mass of vehicle i, and $u_i(t,x)$ is the control input. A common control strategy designates a lead vehicle, while each follower adjusts its motion based on relative distance and velocity measurements from its immediate neighbors (see Figure 1.1). While intuitive and easy to implement, this approach suffers from poor scalability. Small disturbances at the front of the formation may be slightly amplified by each subsequent vehicle. As these deviations accumulate downstream, the result can be large transient errors at the tail of the formation. This effect—where initially small perturbations grow with network size—is known as *string instability* [Stüdli et al., 2017a; Feng et al., 2019; Seiler et al., 2004; Abolfazli et al., 2023]. The phenomenon is illustrated in Figure 1.2, where increasing the number of vehicles leads to drastically increased deviations in transient response.

Figure 1.1 A vehicle platoon where each vehicle only measures distances to its nearest predecessor neighbor.

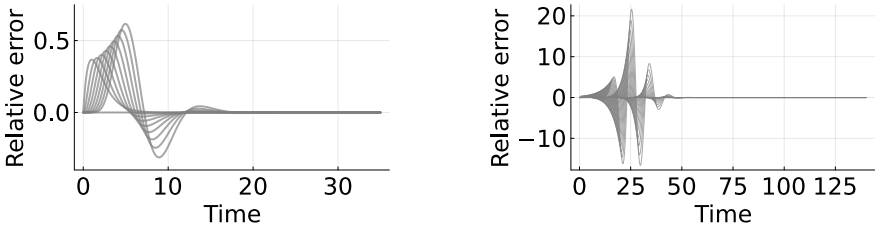

(a) $N = 10$ vehicles. **(b)** $N = 40$ vehicles.

Figure 1.2 Transient response for conventional consensus with varying number of vehicles N. As formation size increases, transient deviations grow rapidly, illustrating string instability.

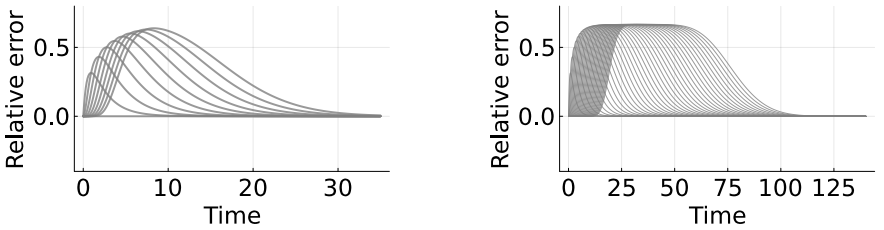

(a) $N = 10$ vehicles. **(b)** $N = 40$ vehicles.

Figure 1.3 Transient response for serial consensus with varying number of vehicles N. Serial consensus maintains bounded performance independent of formation size.

Further complications arise in circular formations, where traditional consensus protocols can become unstable as the number of agents increases [Tegling et al., 2023; Stüdli et al., 2017b]. These observations underscore fundamental performance limitations in conventional distributed control strategies.

Motivated by these challenges, this thesis proposes a compositional design approach termed *compositional consensus*. This framework constructs high-order coordination protocols by serially connecting stable first-order consensus systems. In the linear time-invariant setting, this leads to the closed-loop dynamics:

$$\left(\frac{\mathrm{d}}{\mathrm{d}t} + L_2\right)\left(\frac{\mathrm{d}}{\mathrm{d}t} + L_1\right) x(t) = u_{\mathrm{ref}}(t), \tag{1.2}$$

where L_1 and L_2 are graph Laplacians encoding localized feedback. This architecture is referred to as the *serial consensus* protocol.

While serial consensus may require additional relative measurements—such as from neighbors' neighbors—it fundamentally mitigates the performance degradation seen in traditional consensus designs. In particular, it guarantees that transient errors remain bounded independently of the number of agents, as illustrated in Figure 1.3.

Allowing just one additional relative measurement—for instance, accessing the state of a neighbor's neighbor—yields substantial performance improvements. In a

Figure 1.4 Extended vehicle platoon architecture: each vehicle measures the distance to its immediate predecessor and message-passes this measurement to its follower.

platoon, this corresponds to each vehicle measuring relative positions to both its immediate predecessor and the next vehicle ahead. This can be implemented using a signaling layer, as shown in Figure 1.4. This minor modification enables string-stable behavior while preserving low communication and sensing overhead.

The scalability and robustness of serial consensus are captured by the bound

$$\sup_{t\geq 0} \left\| \begin{bmatrix} e_{\text{pos}}(t) \\ e_{\text{vel}}(t) \end{bmatrix} \right\|_\infty \leq \alpha \left\| \begin{bmatrix} e_{\text{pos}}(0) \\ e_{\text{vel}}(0) \end{bmatrix} \right\|_\infty, \tag{1.3}$$

where $e_{\text{pos}}(t)$ and $e_{\text{vel}}(t)$ denote position and velocity errors, respectively, and the constant α is independent of the number of agents.

This thesis addresses core challenges in coordinating large-scale networked systems under locality constraints—namely, how to design control strategies that remain robust, stable, and performance-guaranteed as the number of agents increases. By explicitly quantifying transient behavior and proposing structurally scalable designs, the thesis contributes both theoretical insights and practical tools for analyzing and controlling distributed multi-agent systems.

Thesis Structure The thesis is organized as follows:

- **Chapter 2** introduces mathematical preliminaries, including graph theory, consensus protocols, and pseudospectral tools for transient analysis.
- **Chapter 3** presents a literature review and contextualizes the thesis's contributions. The proposed coordination protocols are illustrated through their application to multi-agent coordination problems.
- **Chapter 4** summarizes the main results and provides a detailed account of the five appended research articles.
- **Chapter 5** concludes with a discussion of implications, limitations, and directions for future work.

The core of the thesis consists of five appended papers, which are included in full after the introductory chapters.

2

Mathematical Preliminaries

This chapter establishes the mathematical foundations for analyzing and controlling networked multi-agent systems. It introduces key modeling frameworks, distributed control principles, performance norms, and analytical tools for assessing stability and transient behavior. These concepts form the basis for understanding the coordination protocols and performance limitations addressed in later chapters.

2.1 Modeling of Networked Systems

Many real-world systems consist of a large number of interconnected subsystems or agents. Examples include vehicle platoons, drone formations, robotic networks, and power grids. When the individual dynamics are structurally similar, describing the overall system using a compact collective model is advantageous. A common abstraction for vehicular formations is the double-integrator model:

$$M\ddot{x}(t) = u(x,t),$$

where $x(t) \in \mathbb{R}^N$ denotes the positions of N agents, M is a diagonal mass matrix, and $u(x,t)$ is the vector of control inputs.

A typical n^{th}-order linear time-invariant networked system can be described by

$$A_n x^{(n)}(t) + A_{n-1} x^{(n-1)}(t) + \cdots + A_1 \dot{x}(t) + A_0 x(t) = u(t),$$

where $A_k \in \mathbb{R}^{N\times N}$ represent the interconnection structure. Such representations are applicable for modeling systems linearized about equilibrium points. This system can be represented in controllable canonical form as

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{bmatrix} x(t) \\ \vdots \\ x^{(n-2)}(t) \\ x^{(n-1)}(t) \end{bmatrix} = \underbrace{\begin{bmatrix} & I & & \\ & & \ddots & \\ & & & I \\ -A_0 & -A_1 & \cdots & -A_n \end{bmatrix}}_{\mathcal{A}} \begin{bmatrix} x(t) \\ \dot{x}(t) \\ \vdots \\ x^{(n-1)}(t) \end{bmatrix} + \begin{bmatrix} 0 \\ \vdots \\ 0 \\ I \end{bmatrix} u(t), \tag{2.1}$$

where $\mathcal{A}$ is a block companion matrix representing the internal system dynamics.

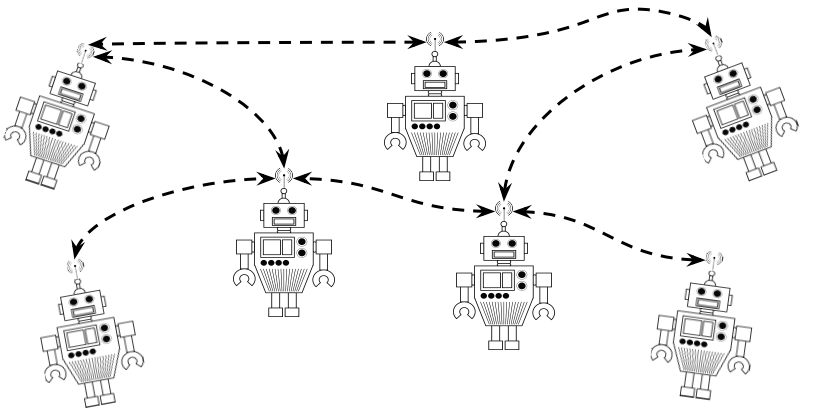

Figure 2.1 A networked system in which the edges indicate communication or coupling between agents.

2.2 Distributed Control for Networked Systems

Large-scale interconnected systems require distributed control strategies to address scalability and communication constraints. A typical linear distributed controller can be expressed as

$$u_i(t) = u_{i,\mathrm{ref}}(t) - \sum_{j \in \mathcal{N}_i} A_{i,j} x_j(t),$$

where $\mathcal{N}_i$ denotes the neighbors of agent i, defined via a communication graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, with nodes $\mathcal{V} = \{1, \ldots, N\}$ and edges $\mathcal{E}$. An illustration is shown in Figure 2.1.

A common design constraint is that control actions depend only on relative state measurements, which is encoded by the condition

$$\sum_{j \in \mathcal{N}_i} A_{i,j} = 0, \quad \forall i.$$

When this condition is satisfied and all off-diagonal elements are non-positive ($A_{ij} \leq 0$ for $i \neq j$), the matrix A is referred to as a *graph Laplacian*, typically denoted L.

The following is a classical result for consensus protocols with constant graph Laplacians.

Lemma 1—[Lin et al., 2005; Ren et al., 2005]
The system $\dot{x} = -Lx$ achieves consensus (i.e., $|x_i(t) - x_j(t)| \to 0$ for all i, j) if and only if the graph associated with L contains a directed spanning tree.

The condition ensures that at least one agent's state influences the entire network, enabling convergence to a common value.

Time-Varying and Delayed Consensus

The above result can be extended to systems with time-varying interconnections and communication delays. The following definition captures the relevant graph-theoretic structure:

Definition 1—δ-digraph [Moreau, 2004]
Let M be a Metzler matrix with zero row sums. The associated δ-digraph is the directed graph with node set $\{1,\dots,N\}$, where an edge from i to j exists if and only if $M_{i,j} \geq \delta$.

Consider a time-delayed consensus protocol of the form

$$\dot{x}_i(t) = -\sum_{j\in\mathcal{N}_i} L_{i,j}(t)\,(x_i(t) - x_j(t-\tau))\,.$$

Define the time-integrated Laplacian as

$$M(t) = \int_t^{t+T} -L(\tau)\,d\tau.$$

If there exist constants $T > 0$, $\delta > 0$, and a fixed node k that serves as the root of a directed spanning tree in the δ-digraph of $M(t)$ for all t, then the system achieves uniform exponential consensus stability.

This result also extends to the case of heterogeneous and time-varying delays $\tau_{i,j}(t) \leq \tau_{\max}$, as shown in [Lu and Liu, 2017]. The analysis is based on bounding the spread between maximum and minimum state values:

$$D^+ \max_{\sigma\in[t-\tau_{\max},t]} x_i(\sigma) - D^+ \min_{\sigma\in[t-\tau_{\max},t]} x_i(\sigma) \leq 0,$$

where D^+ denotes the upper Dini derivative.

In the special case of symmetric feedback and constant delays, the consensus protocol

$$\dot{x}_i(t) = -\sum_{j\in\mathcal{N}_i} L_{i,j}\,(x_i(t-\tau) - x_j(t-\tau))$$

is stable if $\tau \in [0, \tau^*)$, where $\tau^* = \pi/(2\lambda_{\max}(L))$, as shown in [Olfati-Saber and Murray, 2004].

2.3 System Norms and Performance Metrics

A variety of norms are used to quantify system stability, robustness, and performance. These provide meaningful measures of magnitude for both state vectors and system responses.

Norms

This work makes use of both vector norms and their corresponding induced matrix norms. For a vector $z \in \mathbb{C}^N$, the most commonly used norms include:

- **1-norm:** $\|z\|_1 = \sum_{i=1}^{N} |z_i|$,
- **2-norm (Euclidean):** $\|z\|_2 = \sqrt{\sum_{i=1}^{N} |z_i|^2}$,
- **Infinity norm:** $\|z\|_\infty = \max_{i \in \{1,\ldots,N\}} |z_i|$.

The infinity norm is particularly relevant for distributed systems, as it captures the worst-case deviation across agents.

For a matrix $A \in \mathbb{C}^{M \times N}$, the induced norm is defined by

$$\|A\| := \sup_{\|x\|=1} \|Ax\|.$$

The induced matrix norms corresponding to the above vector norms are:

Induced Matrix Norms

- **Max column-sum norm:** $\|A\|_1 = \max_j \sum_i |A_{i,j}|$,
- **Spectral norm:** $\|A\|_2 = \sigma_{\max}(A)$, the largest singular value,
- **Max row-sum norm:** $\|A\|_\infty = \max_i \sum_j |A_{i,j}|$.

These norms will be used to evaluate both internal system behaviors and input-output performance metrics.

2.4 Comparison Functions

Stability properties in nonlinear and time-varying systems are often expressed using comparison functions from the classes $\mathcal{K}$ and $\mathcal{KL}$, which offer compact notations for convergence and boundedness properties.

Definition 2—[Khalil, 2002]
A continuous function $\alpha : [0,a) \to [0,\infty)$ is of class $\mathcal{K}$ if it is strictly increasing and satisfies $\alpha(0) = 0$. If $a = \infty$ and $\alpha(r) \to \infty$ as $r \to \infty$, then $\alpha \in \mathcal{K}_\infty$.

Definition 3—[Khalil, 2002]
A function $\beta : [0,a) \times [0,\infty) \to [0,\infty)$ is of class $\mathcal{KL}$ if, for each fixed s, $\beta(r,s) \in \mathcal{K}$ in r, and for each fixed r, $\beta(r,s)$ is decreasing in s with $\beta(r,s) \to 0$ as $s \to \infty$.

These function classes are used in this thesis to characterize the stability of nonlinear or time-varying consensus protocols. They compactly express notions such as input-to-state stability and global asymptotic stability.

Example 1
A system $\dot{x} = f(t,x)$ is globally asymptotically stable at an equilibrium x^* if there exists $\beta \in \mathcal{K}_\infty\mathcal{L}$ such that

$$\|x(t) - x^*\| \leq \beta(\|x(t_0) - x^*\|, t - t_0), \quad \forall t \geq t_0.$$

2.5 Performance Analysis and Pseudospectra

A matrix A is said to be *normal* if it commutes with its Hermitian adjoint, that is, $AA^{\mathrm{H}} = A^{\mathrm{H}}A$. For linear systems of the form $\dot{x} = Ax$, normality implies that the dynamical behavior can be accurately characterized by the eigenvalues of A. Common examples of normal matrices include symmetric matrices ($A^\top = A$) and skew-symmetric matrices ($A^\top = -A$).

In contrast, *non-normal* matrices—those that do not satisfy the normality condition—can exhibit significant transient growth even when all eigenvalues lie in the left half-plane. This behavior is particularly relevant in the analysis of higher-order consensus systems.

The typical n^{th}-order consensus dynamics (see (2.1)) with $n \geq 2$, where the matrices $A_k = L_k$ are graph Laplacians, are inherently non-normal. This can be seen by inspecting the block system matrix $\mathcal{A}$: the top-left block of $\mathcal{A}\mathcal{A}^{\mathrm{H}}$ equals the identity matrix I, whereas the corresponding block in $\mathcal{A}^{\mathrm{H}}\mathcal{A}$ equals $L_0^\top L_0$, which are not generally equal. This non-normality motivates a more refined study of transient dynamics, as eigenvalue analysis alone may fail to capture the true system behavior.

Non-normality has also been identified as a structural feature of many real-world networks, including biological, social, and engineered systems [Asllani et al., 2018]. In such systems, the directional flow of influence inherent to non-normal interconnection structures can strongly shape transient dynamics. For example, [Baggio et al., 2020] demonstrates how non-normality can be exploited to selectively amplify desired signals in noisy communication networks, in contrast to normal (symmetric) networks, which suffer from poor signal-to-noise ratios. These effects are closely tied to the shape of the system's pseudospectra—a key tool for quantifying transient amplification in non-normal systems. A comprehensive treatment of pseudospectral theory is provided in [Trefethen and Embree, 2005].

Pseudospectra offer a more refined tool than eigenvalues in predicting transient growth. The ε-pseudospectrum of a matrix A is defined as

$$\sigma_\varepsilon(A) = \left\{ s \in \mathbb{C} \;\middle|\; \|(sI - A)^{-1}\| > \frac{1}{\varepsilon} \right\},$$

which captures regions in the complex plane where the resolvent norm is large.

The Kreiss bound [Trefethen and Embree, 2005, Theorem 15.4] provides a lower bound on the worst-case transient growth:

$$\|e^{tA}\| \geq \frac{\alpha_\varepsilon(A)}{\varepsilon} \quad \forall \varepsilon > 0,$$

where $\alpha_\varepsilon(A) = \sup_{s\in\sigma_\varepsilon(A)} \mathrm{Re}(s)$ is the *pseudospectral abscissa*.

An upper bound can also be constructed using the pseudospectrum [Trefethen and Embree, 2005, Theorem 15.2]:

$$\|e^{tA}\| \leq \frac{L_\varepsilon\, e^{t\alpha_\varepsilon(A)}}{2\pi\varepsilon} \quad \forall t \geq 0,$$

where L_ε denotes the arc length of the boundary of the convex hull of $\sigma_\varepsilon(A)$. These bounds highlight how even spectrally stable systems can exhibit significant transient growth due to non-normality.

2.6 Transient Performance and Input-Output Analysis

In control systems, performance is typically best assessed in terms of input-output behavior. This thesis adopts a characterization based on input-output pseudospectra.

Consider the linear time-invariant system:

$$\dot{x}(t) = Ax(t) + Bu(t), \qquad y(t) = Cx(t),$$

with transfer matrix $G(s) = C(sI-A)^{-1}B$. In Paper V, a corresponding input-output ε-pseudospectrum is introduced, formally defined as

$$\sigma_\varepsilon(A,B,C) = \left\{ s \in \mathbb{C} \;\middle|\; \|C(sI-A)^{-1}B\| > \frac{1}{\varepsilon} \right\}.$$

Using this, the worst-case transient response—due either to initial conditions or impulse inputs—can be lower bounded as

$$\sup_{t\geq 0} \|Ce^{tA}B\| \geq \frac{\alpha_\varepsilon(A,B,C)}{\varepsilon}, \quad \forall \varepsilon > 0,$$

where $\alpha_\varepsilon(A,B,C)$, the *input-output pseudospectral abscissa*, is defined as

$$\alpha_\varepsilon(A,B,C) = \sup_{s\in\sigma_\varepsilon(A,B,C)} \mathrm{Re}(s).$$

The following upper bound also holds:

$$\sup_{t\geq 0} \|Ce^{tA}B\| \leq \frac{L_\varepsilon(A,B,C)\, e^{t\alpha_\varepsilon(A,B,C)}}{2\pi\varepsilon},$$

where the constant $L_\varepsilon(A,B,C)$ denotes the arc length of the boundary of the convex hull of $\sigma_\varepsilon(A,B,C)$.

These bounds, along with additional characterizations of the transient behavior of $\|y(t)\|$, are developed in detail in Paper V.

3

Multi-agent Coordination

This chapter motivates the coordination framework developed in this thesis by examining the evolution of distributed control strategies in networked systems. A central theme is how locality constraints—where agents rely solely on relative information from nearby neighbors—impact stability, performance, and scalability.

To contextualize the thesis's contributions, the chapter begins with a review of historical developments in multi-agent coordination, from early consensus protocols to recent work on the formation control of networked systems with high-order dynamics.

3.1 Development of Coordination Protocols

The study of coordination in multi-agent systems has a long and multidisciplinary history, with foundational contributions arising in fields such as transportation systems, opinion dynamics, and biological networks. This section surveys the literature underpinning the coordination strategies explored in this chapter, focusing in particular on second- and high-order consensus protocols, vehicle formations, and string stability.

Introducing the Coordination Problem

Early work on vehicle platooning motivated some of the first formal treatments of decentralized coordination. Seminal contributions include [Levine and Athans, 1966; Melzer and Kuo, 1971; Chu, 1974b; Chu, 1974a], where second-order consensus protocols were proposed to synchronize both position and velocity in chains of vehicles.

In parallel, coordination behavior was studied in other domains. The Kuramoto model [Kuramoto, 1975] addressed nonlinear synchronization in oscillator networks, while DeGroot's model [DeGroot, 1974] introduced a discrete-time opinion dynamics framework based on local averaging. These studies gave rise to a rich literature on bounded-confidence dynamics and social learning, including the Friedkin–Johnsen model [Friedkin and Johnsen, 1990], which remains widely cited. Surveys such

as [Strogatz, 2000; Bernardo et al., 2024; Ureña et al., 2019] provide extensive overviews of these developments.

From Stability to Performance

A major conceptual shift in the study of vehicular formations was the realization that classical asymptotic stability does not guarantee satisfactory transient behavior. This led to the introduction of *string stability*, formalized in [Swaroop and Hedrick, 1996], which aims to prevent the amplification of small disturbances along a vehicle string—an effect that can otherwise grow unbounded with network size.

Early analyses showed that string stability cannot generally be achieved using predecessor-relative feedback alone [Seiler et al., 2004]. Message-passing strategies were proposed in [Barooah and Hespanha, 2005] as one approach to mitigating error amplification. The concept of string stability has since evolved, and its multiple definitions and interpretations are reviewed in the surveys [Stüdli et al., 2017a; Feng et al., 2019].

Experimental studies have reinforced the practical importance of these issues. For instance, [Gunter et al., 2021] demonstrated that modern commercial adaptive cruise controllers do not ensure string-stable behavior. Conversely, [Stern et al., 2018] showed that inserting a small number of autonomous vehicles into a human-driven formation can drastically reduce disturbance propagation and improve fuel efficiency.

Structural Properties of Feedback

Subsequent research emphasized the role of feedback architecture and graph topology in achieving scalable coordination. It was observed in [Barooah et al., 2009; Herman et al., 2017b; Herman et al., 2017a] that combining asymmetric velocity feedback with symmetric position feedback can yield significantly improved performance. This structure also permits the use of passivity theory to establish scalability, as used in [Herman et al., 2017a] to derive bounds of type $\mathcal{L}_2 \to \mathcal{L}_2$.

Other approaches have examined *coherence*—quantifying the global deviation from a mean trajectory—and scalable input-to-state stability (sISS) [Bamieh et al., 2012; Pates et al., 2017; Patterson and Bamieh, 2010; Besselink and Knorn, 2018].

High-Order Consensus

Higher-order consensus protocols were first formalized in [Ren et al., 2006; Ren et al., 2007], which introduced the general form

$$x^{(n)} = -r_{n-1}Lx^{(n-1)} - \cdots - r_1 L\dot{x} - r_0 Lx.$$

Subsequent work explored extensions, including average consensus with absolute feedback [Rezaee and Abdollahi, 2015], delay robustness [Yu et al., 2011; Trindade et al., 2025; Tian and Zhang, 2012], and optimal control-based designs using LQR formulations [Radmanesh et al., 2017; Trindade et al., 2024; Li et al., 2024]. Surveys

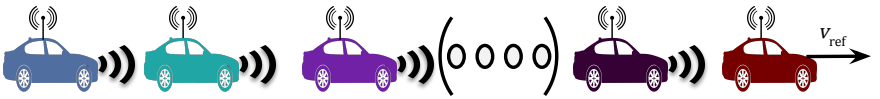

Figure 3.1 Vehicle formation where each car uses onboard sensors to coordinate with its predecessor.

such as [Huang et al., 2014; Amirkhani and Barshooi, 2022] offer broad coverage of this literature.

However, high-order protocols also suffer from *scale fragility*. As shown in [Tegling et al., 2023], if the smallest eigenvalues of the Laplacian decrease with network size, the system may become unstable unless the feedback gains are scaled appropriately. In the case of symmetric banded Laplacians, this decay can be characterized using the Cheeger constant and Buser's inequality [Buser, 1982]. For directed graphs, which are not constrained by this bound, string instability can arise due to poor damping properties [Middleton and Braslavsky, 2010].

Toward Scalable Coordination

The literature reviewed above reveals fundamental limitations in traditional coordination protocols. While early strategies based on symmetric feedback and simple topologies provide asymptotic guarantees, they often exhibit poor transient performance in large-scale formations. In particular, string instability has been shown to be unavoidable in some types of directed networks [Middleton and Braslavsky, 2010].

Several recent works have explored improvements through asymmetric feedback structures or alternative topologies. However, achieving scalable performance in high-order systems with only local, relative feedback remains a largely open problem.

This motivates a different approach: rather than tuning or modifying conventional protocols, the subsequent sections develop a compositional design framework that inherently ensures robustness and scalability. The starting point is a closer examination of classical consensus strategies and their behavior under increasingly realistic models of vehicle formations.

3.2 First-Order Consensus

This section introduces the continuous-time consensus protocol and the modeling framework used throughout the chapter. In particular, several graph Laplacians are defined to represent local relative feedback structures in vehicle formations.

Modeling Setup and Graph Laplacians

This chapter focuses on the control of vehicle formations—illustrated in Figure 3.1—where increasingly sophisticated protocols are introduced to address coordination under locality constraints.

The examples throughout this chapter follow a leader-follower architecture. The feedback structures are encoded using adjacency matrices, representing undirected paths, directed paths, and directed cycles—all configured with a designated leader:

$$[W_{\text{undir}-\text{path}}]_{i,j} = \begin{cases} 1 & \text{if } i \geq 2 \text{ and } i = j \pm 1 \\ 0 & \text{otherwise} \end{cases}, \tag{3.1}$$

$$[W_{\text{ahead}-\text{path}}]_{i,j} = \begin{cases} 1 & \text{if } i \geq 2 \text{ and } i = j + 1 \\ 0 & \text{otherwise} \end{cases}, \tag{3.2}$$

$$[W_{\text{ahead}-\text{cycle}}]_{i,j} = \begin{cases} 1 & \text{if } i = 2 \text{ and } j \in \{1, N\} \\ 1 & \text{if } i \geq 2 \text{ and } i = j + 1 \\ 0 & \text{otherwise} \end{cases}. \tag{3.3}$$

The corresponding graph Laplacians are defined in the standard way:

$$L_{i,j} = \begin{cases} -W_{i,j} & \text{if } i \neq j \\ \sum_{j=1}^{N} W_{i,j} & \text{if } i = j \end{cases}.$$

For example, the Laplacian $L_{\text{ahead}-\text{path}}$ encodes a feedback structure in which each agent observes only its immediate predecessor—i.e., ahead-looking coordination, as depicted in Figure 3.1. With these definitions in place, the consensus protocol can be formalized.

The Consensus System

A natural starting point in vehicle formation control is to ensure that agents converge to a formation with fixed inter-vehicular distances, denoted $d_{\text{ref}} \in \mathbb{R}^N$. In scenarios where vehicles operate at low speeds and can adjust velocities rapidly, the dynamics of each agent can be approximated by the simplified model:

$$\dot{x}_i = u_i(x,t), \tag{3.4}$$

where x_i is the position of agent i, and u_i is the control input determining its velocity.

Remark 1
The model in (3.4) assumes infinite acceleration, which is physically unrealistic. However, it can be justified by time-scale separation, where velocity tracking dynamics are assumed to be much faster than position dynamics.

Assuming that each agent knows its designated offset in the formation, a simple relative feedback law can be implemented:

$$u_i(x,t) = -\sum_j W_{i,j} \left((x_i - d_{\text{ref},i}) - (x_j - d_{\text{ref},j}) \right). \tag{3.5}$$

In vector form, this becomes:

$$u(x,t) = -L(x - d_{\text{ref}}),$$

where L is the graph Laplacian of the communication topology. Introducing the translated state $\tilde{x} = x - d_{\text{ref}}$, the closed-loop dynamics reduce to

$$\dot{\tilde{x}} = -L\tilde{x},$$

which corresponds to the classical consensus protocol, as studied in [Olfati-Saber and Murray, 2004].

The control objective is for each agent to reach its designated position relative to the formation. This is equivalent to:

$$\lim_{t\to\infty} |(x_i(t) - d_{\text{ref},i}) - (x_j(t) - d_{\text{ref},j})| = 0, \quad \forall i, j,$$

which implies consensus in the shifted coordinates $\tilde{x}(t)$.

Definition 1—Consensus
A solution $x(t)$ is said to achieve consensus if

$$\lim_{t\to\infty} |x_i(t) - x_j(t)| = 0, \quad \forall i, j.$$

In the context of vehicle formations, consensus in the shifted state $\tilde{x}(t)$ means that all vehicles maintain the desired inter-vehicular spacing defined by d_{ref}. The formation converges asymptotically to a trajectory of the form

$$x(t) = d_{\text{ref}} + a(t)\mathbf{1},$$

for some scalar function $a(t)$, representing collective motion.

A large body of literature addresses the stability and robustness of consensus protocols. These results extend to settings with time-varying topologies, nonlinearities, input saturation, and communication delays; see, e.g., [Moreau, 2004; Lin et al., 2005; Lyu et al., 2016; Olfati-Saber and Murray, 2004], as well as the summary in Section 2.2. Collectively, these results establish that consensus can be robustly achieved under a broad range of practical conditions.

For instance, the Laplacians

$$L_{\text{ahead}-\text{path}}, \quad L_{\text{undir}-\text{path}}, \quad L_{\text{ahead}-\text{cycle}}$$

each correspond to a graph with a directed spanning tree, and therefore yield consensus in the translated state $\tilde{x}$ using the protocol (3.5).

Furthermore, robustness can be shown for formations with heterogeneous agents. Suppose a control law is designed for a communication graph that contains a directed spanning tree. If each vehicle uses a time-varying gain $k_i(t)$, the dynamics become

$$\dot{x}_i = k_i(t)u_i(x,t).$$

Then, as shown in [Moreau, 2004], consensus is preserved as long as $k_i(t) \geq k_{\min} > 0$ for all i and t.

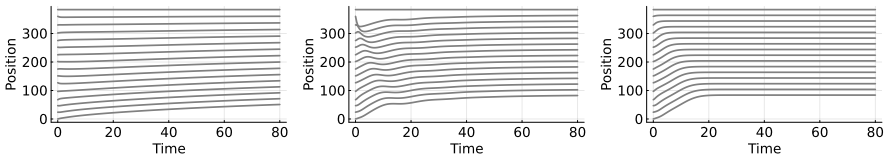

(a) Undirected path graph **(b)** Directed cycle graph **(c)** Directed path graph

Figure 3.2 Simulation of first-order consensus under different graph topologies. All structures lead to convergence to the desired formation.

Simulation Study

Simulation Setup To illustrate these concepts, consider the following second-order vehicle model:

$$m_i\ddot{x}_i = -p_i(\dot{x}_i - v_{\text{ref},i}),$$

where m_i and p_i denote the mass and control gain, respectively. The reference velocity is computed using the controller (3.5), i.e.,

$$v_{\text{ref}} = -L\tilde{x}.$$

Simulation Results Figure 3.2 shows the evolution of vehicle positions under three different feedback structures. In all cases, the formation converges to the desired configuration after a transient period.

3.3 Second-Order Consensus

Consider a collection of agents with second-order integrator dynamics, modeled by

$$\ddot{x}_i = u_i(x,t),$$

or, in compact vector form, $\ddot{x} = u(t,x)$. This simple model was already considered in early studies of large vehicle strings, such as [Levine and Athans, 1966; Melzer and Kuo, 1971; Chu, 1974b; Chu, 1974a]. In this setting, the control objective is to coordinate the agents in both position and velocity.

Definition 2—Second-order consensus
A solution $x(t)$ is said to achieve second-order consensus if

$$\lim_{t\to\infty} |x_i(t) - x_j(t)| = 0 \quad \textit{and} \quad \lim_{t\to\infty} |\dot{x}_i(t) - \dot{x}_j(t)| = 0, \quad \forall i,j.$$

A frequently studied control law that relies only on relative measurements takes the form:

$$u_i(x,t) = -\sum_{j\in\mathcal{N}_i} [W_{\text{pos}}]_{i,j}\left((x_i - d_{\text{ref},i}) - (x_j - d_{\text{ref},j})\right) + [W_{\text{vel}}]_{i,j}(\dot{x}_i - \dot{x}_j),$$

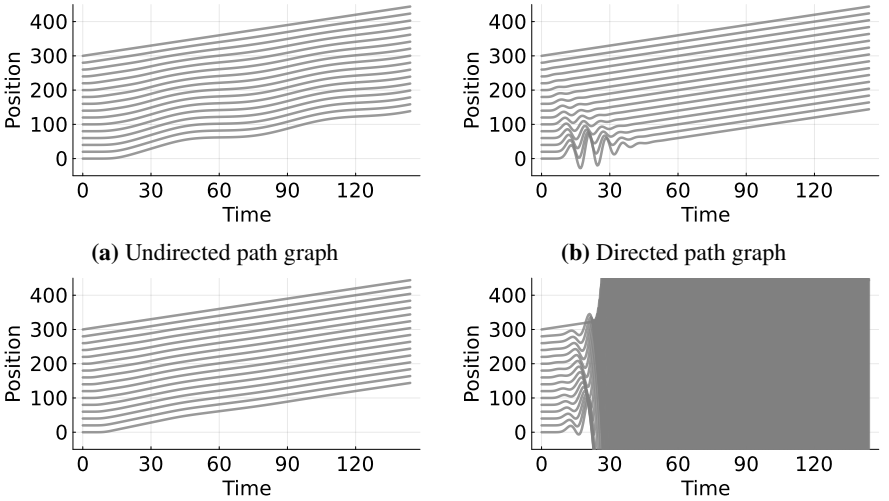

(a) Undirected path graph

(b) Directed path graph

(c) Undirected and directed path graph

(d) Directed cycle graph

Figure 3.3 Second-order consensus simulations for four vehicle formations with different network structures. Observed behaviors include macroscopic wave propagation (a), string instability (b), improved performance through mixed topologies (c), and outright instability (d).

or, in vector form using graph Laplacians:

$$u(t,x) = -L_{\text{vel}}\dot{x} - L_{\text{pos}}(x - d_{\text{ref}}). \tag{3.6}$$

Designing a protocol that achieves second-order consensus is straightforward when symmetric Laplacians are used for both position and velocity feedback:

$$L_{\text{pos}} = r_{\text{pos}}L, \qquad L_{\text{vel}} = r_{\text{vel}}L.$$

However, deviations from this structure—such as directed graphs or non-uniform weighting—can lead to undesirable behaviors as the network grows. These include string instability [Seiler et al., 2004] and, in some cases, instability [Stüdli et al., 2017b]. This is illustrated in the following example.

Example 1—Second-order consensus simulation

Four second-order consensus protocols are simulated based on a leader-follower architecture, defined by (3.1)–(3.3), using the following graph structures:

- ***Undirected path:*** $L_{\text{pos}} = L_{\text{vel}} = L_{\text{undir-path}}$
- ***Directed path:*** $L_{\text{pos}} = L_{\text{vel}} = L_{\text{ahead-path}}$

- ***Undirected and directed path:*** $L_{\text{pos}} = L_{\text{undir-path}}$, $L_{\text{vel}} = L_{\text{ahead-path}}$
- ***Directed cycle:*** $L_{\text{pos}} = L_{\text{vel}} = L_{\text{ahead-cycle}}$

The leader (agent 1) follows a constant-velocity reference trajectory, $\dot{x}_1(t) = v_0$, while the other agents apply the protocol (3.6). The simulation results in Figure 3.3 show qualitatively distinct transient behaviors:

- ***Symmetric path:*** *(Figure 3.3a): long-lasting macroscopic wave patterns.*
- ***Directed path:*** *(Figure 3.3b): string instability with downstream error amplification.*
- ***Undirected and directed path:*** *(Figure 3.3c): improved performance, consistent with [Herman et al., 2017a].*
- ***Directed cycle:*** *(Figure 3.3d): instability for sufficiently large networks [Stüdli et al., 2017b; Tegling et al., 2023].*

These transient behaviors are further explained through input-output ε-pseudospectra. The transfer function from the leader's velocity to the local position error is given by:

$$G_{e_{\text{pos}},\dot{x}_0}(s) = L_{\text{ahead-path}}(s^2 I + sL_{\text{vel}} + L_{\text{pos}})^{-1},$$

where $e_{\text{pos}}(t) = L_{\text{ahead-path}}(x(t) - d_{\text{ref}})$. As discussed in Paper V, the worst-case transient response $\|e_{\text{pos}}(t)\|_\infty$ can be bounded by the input-output pseudospectrum of this system:

$$\sigma_\varepsilon = \left\{ s \in \mathbb{C} \mid \|G_{e_{\text{pos}},\dot{x}_0}(s)\|_\infty > \frac{1}{\varepsilon} \right\}.$$

Figure 3.4 shows the level curves of σ_ε for $N = 50$, along with the system poles. In the case of the directed cycle graph (Figure 3.4d), poles lie in the right-half plane, indicating instability. For the directed path topology, the system remains asymptotically stable, but the pseudospectrum extends deep into the right-half plane for small ε. For example,

$$\|G_{e_{\text{pos}},\dot{x}_0}(0.1 + 0.9\text{i})\|_\infty \approx 3.0 \cdot 10^5,$$

suggesting extreme transient amplification. In particular, Theorem 1 in Paper V establishes that for some initial condition satisfying $\|\dot{x}(0)\|_\infty \leq 1$, the local error will amplify to $\|e_{\text{pos}}(t)\|_\infty \geq 3 \cdot 10^4$ for some $t \geq 0$.

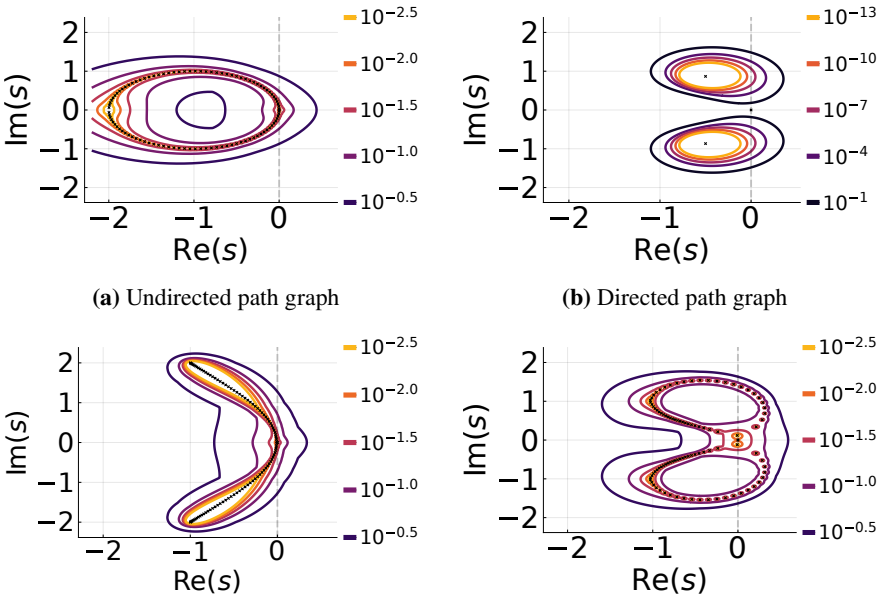

(a) Undirected path graph

(b) Directed path graph

(c) Undirected and directed path graph

(d) Directed cycle graph

Figure 3.4 Level curves of the input-output ε-pseudospectra for several topologies. System poles are marked with black '×'. For the directed path topology in (b), the pseudospectra extend deep into the right-half plane, suggesting large transient amplification. The cyclic graph in (d) exhibits unstable poles.

Scalability and Performance Metrics

The example above illustrates the sensitivity of second-order consensus to graph topology and feedback structure. Two related concepts have been widely studied: *coherence* and *string stability*.

Coherence quantifies deviations from an average trajectory in response to disturbances, often in terms of the $\mathcal{H}_2$ norm; see [Patterson and Bamieh, 2010; Bamieh et al., 2012; Pates et al., 2017; Tegling and Sandberg, 2017]. String stability, by contrast, addresses whether disturbances or initial errors amplify along a network. Several formulations of string stability exist; see [Stüdli et al., 2017a; Feng et al., 2019]. While 2-norm-based metrics are analytically convenient, practical applications often demand performance in the ∞-norm. Such measures are less studied [Feintuch and Francis, 2012].

A related notion is scalable input-to-state stability (sISS), as in [Besselink and Knorn, 2018; Silva et al., 2025], which focuses on worst-case bounded disturbances and requires bounds of the form $\sup_{t\geq 0}\|x(t)\|_\infty$. The serial consensus protocol introduced in Paper II and Paper III achieves this stronger ℓ_∞-type string stability—a

key step toward practical, scalable coordination.

3.4 Second-Order Serial Consensus

A central challenge in second-order consensus is designing coordination protocols that guarantee string stability under locality constraints. Two comprehensive surveys on the subject are [Feng et al., 2019; Stüdli et al., 2017a]. Much of the difficulty stems from the requirement that control laws rely only on local and relative feedback.

A widely studied class of linear protocols adhering to this constraint assumes a structure of the form

$$u(x,t) = -L_{\text{vel}}\dot{x} - L_{\text{pos}}x,$$

which clearly is both local and restricted to relative feedback. However, this is not the only way to implement local and relative feedback.

This thesis introduces an alternative coordination architecture called the *serial consensus* framework.

Definition 3—Second-order serial consensus
The second-order serial consensus system is defined by

$$(\frac{\mathrm{d}}{\mathrm{d}t} + L_2)(\frac{\mathrm{d}}{\mathrm{d}t} + L_1)x(t) = u_{\text{ref}}(t), \tag{3.7}$$

where L_1 and L_2 are graph Laplacians.

This system can be viewed as a compositional generalization of the first-order consensus protocol. The key idea is to specify the desired closed-loop dynamics and then identify the controller that achieves these dynamics.

Definition 4—Second-order serial consensus controller
The control law that yields the closed-loop dynamics in (3.7) is given by

$$u(x,t) = -(L_2 + L_1)\dot{x} - L_2L_1x + u_{\text{ref}}(t). \tag{3.8}$$

Like its conventional counterpart, this controller can be implemented using only relative feedback. A convenient criterion for verifying relative feedback implementability is the following:

Lemma 1
A linear control law Ax can be implemented using relative feedback if and only if $A\mathbf{1} = 0$.

Proof. For necessity, consider the i^{th} row of Ax:

$$[Ax]_i = \sum_{j=1}^{N} A_{i,j}x_j = \sum_{j=1}^{N} A_{i,j}x_j - x_i\sum_{j=1}^{N} A_{i,j} = \sum_{j=1}^{N} A_{i,j}(x_j - x_i),$$

Figure 3.5 A vehicle formation implementing a serial consensus controller. Each vehicle measures the relative distance and velocity to the predecessing vehicle and signals the perceived relative distance to its follower.

which holds because $\sum_{j=1}^{N} A_{i,j} = 0$, i.e., $[A\mathbf{1}]_i = 0$. For sufficiency, note that any relative measurement can be expressed as follows: $x_j - x_k = (x_j - x_i) - (x_k - x_i)$. This shows that the i^{th} row can be expressed as

$$[Ax]_i = \sum_{j=1}^{N} A_{i,j}(x_j - x_i).$$

Inserting $x = \mathbf{1}$ shows that $[A\mathbf{1}]_i = 0$. Since i was arbitrary, it follows that $A\mathbf{1} = 0$. □

This condition is satisfied for the serial consensus controller. By definition, both L_1 and L_2 satisfy $L_k\mathbf{1} = 0$, and thus:

$$(L_2 + L_1)\mathbf{1} = 0, \quad L_2L_1\mathbf{1} = L_2(L_1\mathbf{1}) = 0.$$

Therefore, the controller (3.8) is implementable via relative feedback of x and $\dot{x}$.

Stability and Locality

The serial consensus protocol inherits the stability properties of the underlying first-order consensus systems.

Proposition 1
Let $u_{\text{ref}} = 0$ and suppose that L_1 and L_2 are graph Laplacians. Then, the solution $x(t)$ to (3.7) achieves second-order consensus for any initial condition if both L_1 and L_2 contain a directed spanning tree.

Proof. This is one of the consequences of Theorem 1 in Paper 1. □

The structure of the controller also respects the locality constraints: if L_1 and L_2 are defined over a graph with adjacency matrix W, then: $L_2 + L_1$ uses 1-hop information; L_2L_1 uses at most 2-hop information. Hence, the controller remains localized, provided that L_1 and L_2 are local. Figure 3.5 illustrates how the second-order serial consensus controller can be implemented using a signaling layer.

Scalable Performance

Beyond stability, the serial consensus controller guarantees scalable performance. Consider the perturbed system:

$$\ddot{x} = -(p_2 + p_1)L\dot{x} - p_2 p_1 L^2 x + Lw_0,$$

where $w_0 \in \mathbb{R}^N$ is a constant load disturbance and L is a graph Laplacian. Define the local position and velocity errors:

$$e_{\mathrm{p}}(t) = L(x(t) - d_{\mathrm{ref}}), \quad e_{\mathrm{v}}(t) = \dot{x}(t) - v_{\mathrm{ref}}\mathbf{1}.$$

Then the following bound holds:

$$\left\| \begin{bmatrix} e_{\mathrm{p}}(t) \\ e_{\mathrm{v}}(t) \end{bmatrix} \right\|_\infty \leq \frac{p_1 + p_2 + 2\max\{1, p_1 p_2\}}{|p_1 - p_2|} \left\| \begin{bmatrix} e_{\mathrm{p}}(0) \\ e_{\mathrm{v}}(0) \end{bmatrix} \right\|_\infty + \alpha_w(p_1, p_2)\|w_0\|_\infty,$$

where $\alpha_w(p_1, p_2)$ is a constant.

Proof sketch: Consider first the unperturbed system:

$$\begin{bmatrix} \dot{x} \\ \ddot{x} \end{bmatrix} = \begin{bmatrix} 0 & I \\ -p_1 p_2 L^2 & -(p_1 + p_2)L \end{bmatrix} \begin{bmatrix} x \\ \dot{x} \end{bmatrix}.$$

Introducing the error variable $e = Lx$ leads to the transformed system:

$$\begin{bmatrix} \dot{e} \\ \ddot{x} \end{bmatrix} = \left(\underbrace{\begin{bmatrix} 0 & 1 \\ -p_1 p_2 & -(p_1 + p_2) \end{bmatrix}}_{A} \otimes L \right) \begin{bmatrix} e \\ \dot{x} \end{bmatrix}.$$

The matrix A has eigenvalues $-p_1$ and $-p_2$, and is therefore diagonalizable when $p_1 \neq p_2$. In particular, there exists a similarity transformation $A = S^{-1}PS$, where $P = \mathrm{diag}(p_1, p_2)$. This gives the decomposition:

$$-(SPS^{-1} \otimes L) = (S \otimes I)(P \otimes -L)(S^{-1} \otimes I),$$

showing that the system is similar to two decoupled first-order consensus systems.

The performance bound follows from identifying an optimal transformation matrix S that diagonalizes A, where optimal is with respect to the cost $\|S^{-1}\|_\infty \|S\|_\infty$. The constant $\alpha_w(p_1, p_2)$ is derived using similar techniques. Full details, including an algorithm for computing the optimal S, are provided in Paper II.

Simulation Results

The serial consensus protocol is evaluated using the same four interconnection structures as in the conventional second-order simulations:

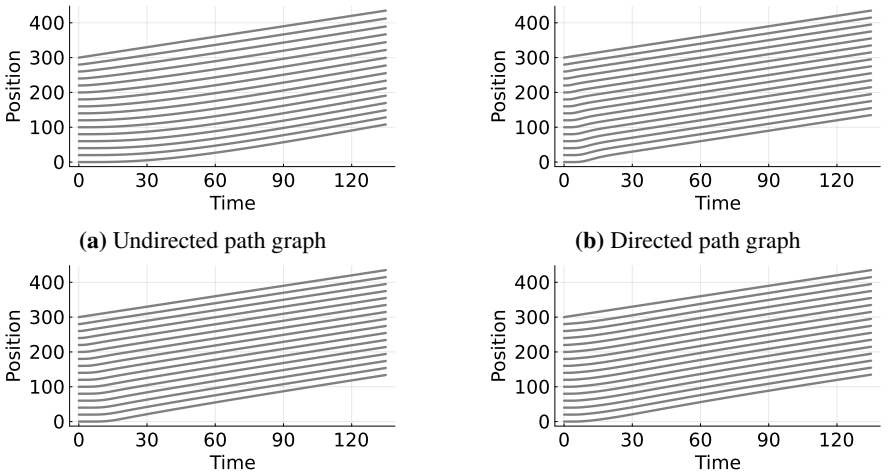

(a) Undirected path graph

(b) Directed path graph

(c) Undirected and directed path graph

(d) Directed cycle graph

Figure 3.6 Serial consensus system under four different graph topologies. The system achieves stable coordination in all cases, with differences primarily in convergence rate.

- **Undirected path:** $L_1 = p_1 L_{\text{undir-path}}, \quad L_2 = p_2 L_{\text{undir-path}}$
- **Directed path:** $L_1 = p_1 L_{\text{ahead-path}}, \quad L_2 = p_2 L_{\text{ahead-path}}$
- **Undirected and directed path:** $L_1 = p_1 L_{\text{ahead-path}}, \quad L_2 = p_2 L_{\text{undir-path}}$
- **Directed cycle:** $L_1 = p_1 L_{\text{ahead-cycle}}, \quad L_2 = p_2 L_{\text{ahead-cycle}}$

The simulations are carried out with $p_1 = 1$, $p_2 = 2$, and a leader agent that maintains constant velocity. As shown in Figure 3.6, the serial consensus protocol achieves convergence with bounded transients across all topologies. In contrast to the conventional protocol, no instability or severe amplification is observed. The main difference between the scenarios lies in the convergence rate.

To further assess the transient behavior, the input-output transfer function from the leader's velocity $\dot{x}_0$ to the local position error $e_{\text{pos}}(t) = L_{\text{ahead-path}}(x(t) - d_{\text{ref}})$ is considered:

$$G_{e_{\text{pos}},\dot{x}_0}(s) = L_{\text{ahead-path}}(sI + L_1)^{-1}(sI + L_2)^{-1}.$$

The corresponding input-output ε-pseudospectra characterize the potential for worst-case transient amplification. Level curves of the pseudospectrum for representative values of ε along with the system poles, are shown in Figure 3.7.

Unlike in the conventional second-order consensus case, none of the pseudospectra for the serial consensus extend far into the right-half plane. This indicates that the worst-case transient growth is modest and consistent with the simulations in Figure 3.6. Even for the directed cycle graph—which was unstable in the conventional protocol—the serial consensus maintains stability and bounded responses.

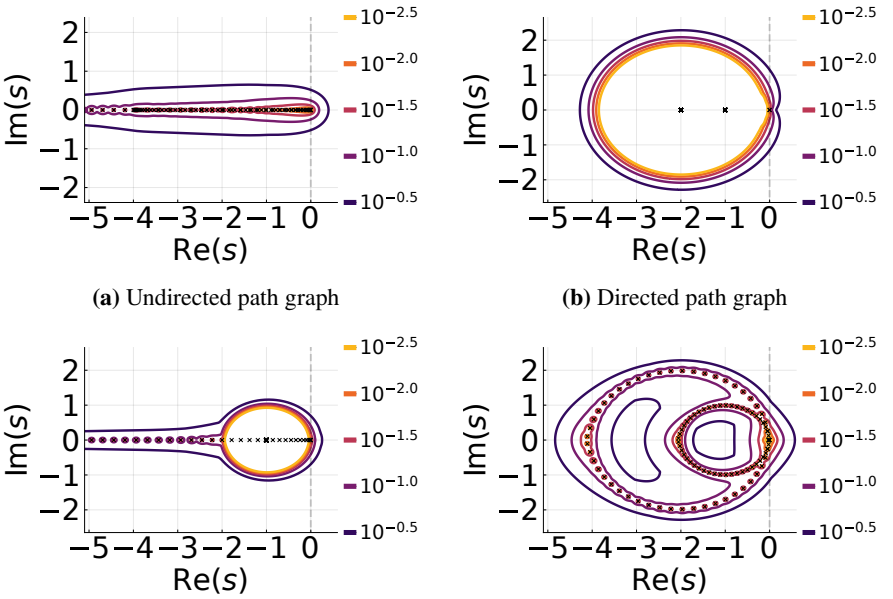

(a) Undirected path graph

(b) Directed path graph

(c) Undirected and directed path graph

(d) Directed cycle graph

Figure 3.7 Level curves of the input-output ε-pseudospectra are plotted for each graph topology, along with system poles (marked by black '×'). In contrast to the conventional second-order protocol, none of the serial consensus configurations exhibit pseudospectra that extend significantly into the right-half plane. This indicate a modest transient amplification of $\|e_{\text{pos}}(t)\|_\infty$.

3.5 Third- and High-Order Consensus

Extending second-order consensus to higher orders introduces new challenges. A natural generalization of the conventional consensus protocol was proposed in [Ren et al., 2006], yielding the following nth-order system:

$$x^{(n)} = -r_{n-1}Lx^{(n-1)} - \cdots - r_1 L\dot{x} - r_0 Lx,$$

where L is a graph Laplacian and $r_k > 0$ are control gains. Numerous early works investigated the properties and stability of this protocol; an overview of its development is given in [Huang et al., 2014].

In vehicle formations, one compelling application of high-order consensus is disturbance rejection. When constant load disturbances act on the system—such as those caused by road slope or aerodynamic drag—integral action is required to eliminate steady-state error. However, due to locality constraints and practical limitations, feedback is often restricted to relative measurements. This naturally

leads to a third-order consensus protocol of the form:

$$\dddot{x} = -r_{\text{vel}} L\ddot{x} - r_{\text{pos}} L\dot{x} - r_{\text{I}} \int_0^t Lx(\tau)\, d\tau.$$

Unlike the second-order case, third- and higher-order conventional consensus protocols suffer from two major limitations:

i) *Scale fragility:* As shown in [Tegling et al., 2023], if the smallest eigenvalues of L approach zero with increasing network size, the gains r_k must be retuned to preserve stability.

ii) *String instability:* Sparse graph Laplacians that help avoid fragility (e.g., directed chains) tend to exhibit severe string instability, where transient error grows exponentially with the number of agents—even in second-order systems [Middleton and Braslavsky, 2010].

The serial consensus framework provides a scalable alternative for high-order coordination. Consider the third-order serial consensus controller:

$$u(x,t) = -(p_1 + p_2 + p_3) L\dot{x} - (p_1 p_2 + p_1 p_3 + p_2 p_3) L^2 (x - d_{\text{ref}}) - p_1 p_2 p_3 Lz$$
$$z(t) = \int_0^t L^2 (x(\tau) - d_{\text{ref}})\, d\tau,$$

where $z(t)$ represents the local integrator state for each agent. This proportional-integral control law ensures stability for any choice of $p_k > 0$, provided that L contains a directed spanning tree.

This controller achieves both disturbance rejection and desired formation maintenance for systems with double-integrator dynamics under constant disturbances. The result generalizes to arbitrary order $n \geq 1$, as shown in Paper III.

Simulation Study

Figure 3.8 presents a simulation comparing second- and third-order serial consensus protocols for a formation of 15 agents subjected to a spatially distributed load disturbance. The third-order proportional-integral (PI) controller eliminates the steady-state error, while the second-order proportional (P) controller converges to a formation with a static offset.

3.6 Nonlinear High-Order Coordination

'All models are wrong..." [Box, 1976], and modeling a physical vehicle as a simple double integrator is no exception. While this abstraction enables elegant control design, transitioning from theory to practice introduces significant challenges. The strong performance of serial consensus in the linear setting naturally raises the

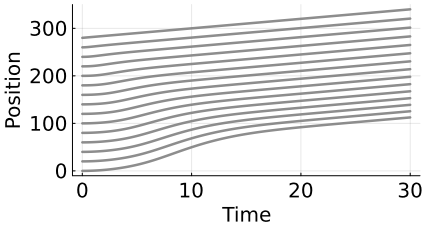

(a) Positions with second-order (P) controller

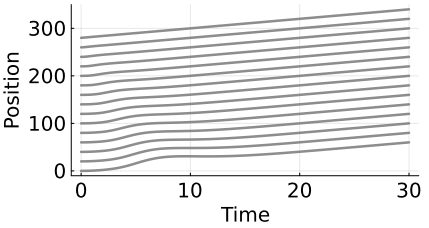

(b) Positions with third-order (PI) controller

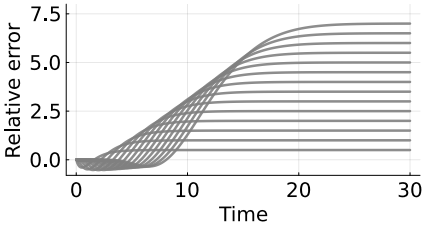

(c) Relative error: $L_{\text{ahead}-\text{path}}(x - d_{\text{ref}})$ with second-order (P) controller

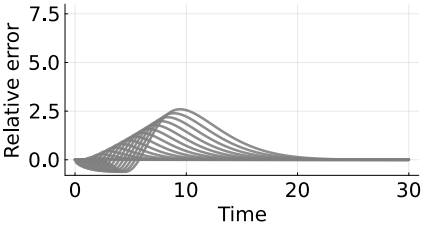

(d) Relative error: $L_{\text{ahead}-\text{path}}(x - d_{\text{ref}})$ with third-order (PI) controller

Figure 3.8 Simulation of a 15-vehicle formation subjected to a spatially distributed load disturbance, using a serial consensus controller. The second-order (P) controller lacks integral action and results in a steady-state offset. In contrast, the third-order (PI) controller incorporates local integrator states, achieving zero steady-state error.

question: Are these benefits an artifact of the idealized assumptions and precisely structured feedback?

This motivates the study of nonlinear and time-varying scenarios. In vehicular formations—as in most control systems—several practical uncertainties must be addressed, including:

- **Actuation delays**: control inputs are inherently discretized in time.
- **Throttle saturations**: vehicles are subject to physical acceleration limits.
- **Heterogeneity**: not all agents share identical dynamics or capabilities.
- **Time-varying dynamics**: road slope and surface conditions can affect acceleration and braking capacity.

Each of these has been considered in the literature to varying degrees; see, e.g., [Chu et al., 2024; Vegamoor et al., 2019; Knorn et al., 2016]. With a small but systematic generalization of the serial consensus protocol, it is possible to design robust controllers under all of the above scenarios.

This leads to the *compositional consensus protocol*, introduced in Paper IV, which generalizes serial consensus to a broad class of (potentially nonlinear and time-varying) coordination dynamics.

Definition 5—Compositional consensus protocol
A system is said to follow a compositional consensus protocol if it satisfies

$$\left(\frac{\mathrm{d}}{\mathrm{d}t}+\mathcal{L}_n\right)\circ\left(\frac{\mathrm{d}}{\mathrm{d}t}+\mathcal{L}_{n-1}\right)\circ\cdots\circ\left(\frac{\mathrm{d}}{\mathrm{d}t}+\mathcal{L}_1\right)x=0, \tag{3.9}$$

where each $\mathcal{L}_k:\mathbb{R}^N\times\mathbb{R}_+\to\mathbb{R}^N$ is a Laplacian-like operator, possibly nonlinear or time-varying.

This structure enables a modular approach to consensus protocol design, where each protocol can be tuned to match the present uncertainty.

Time-Varying Networks To accommodate time-varying topologies, one can let $\mathcal{L}_k(x,t)=L_k(t)x$, where each $L_k(t)$ is a graph Laplacian evolving over time. Provided sufficient smoothness and the conditions in [Moreau, 2005]—namely, sufficient over-time connectivity—high-order consensus is achieved.

Saturation Constraints In settings where actuation is limited, the protocol can incorporate elementwise saturation functions:

$$\mathcal{L}_k(x)=p\,\mathrm{sat}(L_k x),$$

where $\mathrm{sat}(\cdot)$ is the standard saturation function applied elementwise.

Example 2—Saturated second-order consensus
Consider a second-order network with double-integrator dynamics, and let the control law be

$$u(x,t)=p\,\mathrm{sat}\left(L_2\left[\dot{x}+p\,\mathrm{sat}(L_1(x-d_{\mathrm{ref}}))\right]\right)-p\frac{\mathrm{d}}{\mathrm{d}t}\,\mathrm{sat}(L_1(x-d_{\mathrm{ref}})). \tag{3.10}$$

This corresponds to a compositional consensus protocol as in (3.9), with a saturated consensus protocol for each operator $\mathcal{L}_k$. The time derivative of the saturation will satisfy

$$\left[\frac{\mathrm{d}}{\mathrm{d}t}\,\mathrm{sat}(L_1(x-d_{\mathrm{ref}}))\right]_i=[L_1\dot{x}]_i\cdot\mathbf{I}_{(-1,1)}\left([L_1(x-d_{\mathrm{ref}})]_i\right),$$

where $\mathbf{I}_{(-1,1)}(\cdot)$ denotes an indicator function for the linear regime of the saturation.

The closed-loop system has the following form:

$$\begin{bmatrix}\dot{x}\\ \dot{\xi}\end{bmatrix}=\begin{bmatrix}-p\,\mathrm{sat}(L_1(x-d_{\mathrm{ref}}))+\xi\\ -p\,\mathrm{sat}(L_2\xi)\end{bmatrix},$$

where ξ serves as an internal state. It is easy to show

$$\|\xi(t)\|_\infty \le \|\xi(0)\|_\infty, \quad \|\dot{x}(t)\|_\infty \le \|\xi(t)\|_\infty + p, \quad \|\xi(t)\|_\infty \le \|\dot{x}(t)\|_\infty + p,$$

which consequently implies

$$\|\dot{x}(t)\|_\infty \le \|\dot{x}(0)\|_\infty + 2p.$$

This shows that the control input satisfies

$$\|u(x,t)\|_\infty \le p + p\|L_1\|_\infty(\|\dot{x}(0)\|_\infty + 2).$$

Hence, the control signal remains uniformly bounded in time as a function of the initial velocity deviation. Moreover, the bound can be made arbitrarily small by decreasing the gain parameter p. This example is further discussed in Paper IV.

Simulation Results

Figure 3.9 illustrates a simulation of the saturated compositional controller with $\mathcal{L}_1 = \mathcal{L}_2 = p\,\mathrm{sat}(L_{\mathrm{ahead-path}})x$, under two different values of p. The control input plots show that the maximum control effort can be effectively limited by tuning the gain.

The example highlights how the compositional consensus framework extends serial consensus to nonlinear systems. Paper IV discusses further extensions to networks with time-varying dynamics and communication delays.

Summary: Serial and Compositional Consensus

This thesis introduces the serial consensus framework as a scalable alternative to conventional high-order protocols. It guarantees stability and performance under strict locality constraints by composing lower-order consensus systems, as developed in Papers I–III. This design has since been extended to a general class of compositional consensus protocols, enabling nonlinear and time-varying consensus strategies Paper IV.

These developments offer a unifying framework for robust coordination, extending the consensus literature beyond its original linear and time-invariant scope and showing that stable high-order protocols can be systematically constructed—even in networks with heterogeneous agents, saturation effects, or time-varying communication.

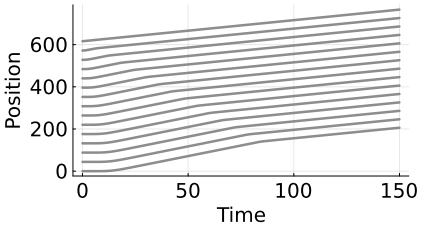

(a) Positions with $p = 1$

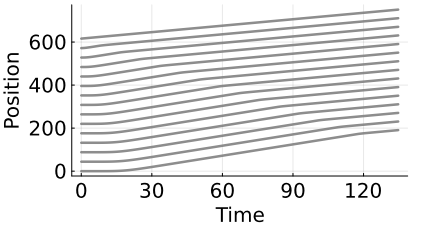

(b) Positions with $p = 0.75$

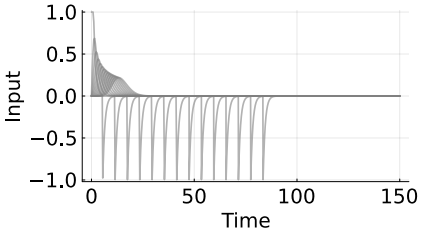

(c) Control inputs with $p = 1$

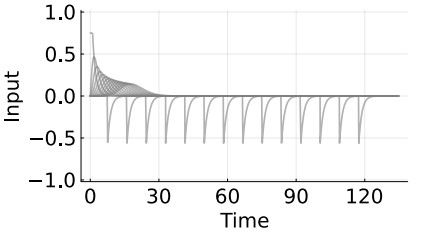

(d) Control inputs with $p = 0.75$

Figure 3.9 Simulation of a vehicle formation using the saturated compositional consensus controller (3.10). Decreasing the gain p reduces the maximum control effort, albeit at the cost of slower convergence.

4

Contribution

This chapter provides a summary of the five appended papers, which together constitute the main scientific contributions of this thesis.

4.1 Included Papers

Paper I

J. Hansson and E. Tegling (2023). "A closed-loop design for scalable high-order consensus". In: *Proc. IEEE Conf. Decis. Control*, pp. 7388–7394

Scientific Summary This paper addresses the problem of coordinating a group of N identical n^{th}-order integrator systems, modeled as

$$\frac{\mathrm{d}^n x(t)}{\mathrm{d}t^n} = u(x,t), \quad x(t) \in \mathbb{R}^N,$$

subject to the constraint that the controller $u(x,t)$ must be linear, local, and rely only on relative state measurements.

A novel control architecture—termed the *serial consensus* controller—is proposed. The controller is designed to produce closed-loop dynamics of the form

$$\left(\prod_{k=1}^{n} sI + L_k\right) X(s) = U_{\text{ref}}(s),$$

where each L_k is a graph Laplacian. Stability of the consensus equilibrium is established under the mild condition that each Laplacian corresponds to a graph containing a directed spanning tree. This resolves a key limitation of conventional high-order consensus, which can become unstable as network size increases unless gains are retuned.

The serial consensus controller preserves the desired structural properties: it is local and implementable using only relative measurements. Specifically, if each L_k

encodes feedback among neighboring agents, then the overall controller requires at most n-hop communication.

Robustness is also analyzed. Using the small-gain theorem, the closed-loop system is shown to tolerate both modeling and feedback uncertainties—providing confidence in its real-world applicability. These results are supported by numerical examples illustrating both stability and robustness under locality constraints.

Paper II

> J. Hansson and E. Tegling (2025a). "Closed-loop design for scalable performance of vehicular formations". *IEEE Trans. Control Netw. Syst.*, pp. 1–10

Scientific Summary This paper investigates a second-order instance of the serial consensus protocol, where the Laplacians used in the coordination law are proportional. The resulting closed-loop system is

$$\ddot{x} = -(p_1 + p_2)L\dot{x} - p_1 p_2 L^2 x + u_{\text{ref}}.$$

The implementation of the term $L^2 x$ is addressed using two alternative strategies:

i) Local relative measurements spanning a two-hop neighborhood.

ii) A message-passing scheme, where each agent transmits its observed relative error to its immediate followers.

A key contribution of the paper is the introduction of a notion of *scalable performance*, defined as bounded amplification from initial conditions to the maximum transient error:

$$\left\| \begin{bmatrix} e_{\text{p}}(t) \\ e_{\text{v}}(t) \end{bmatrix} \right\|_\infty \leq \alpha \left\| \begin{bmatrix} e_{\text{p}}(0) \\ e_{\text{v}}(0) \end{bmatrix} \right\|_\infty .$$

This condition ensures that the transient response remains uniformly bounded as the network scales.

The main result of the paper is a theorem establishing that the serial consensus system satisfies this scalable performance property. Accompanying simulation studies demonstrate that the protocol achieves string-stable behavior—even on directed path graphs—making it particularly well-suited for vehicular formation control.

Paper III

> J. Hansson and E. Tegling (2024). "Performance bounds for multi-vehicle networks with local integrators". *IEEE Control Syst. Lett.* **8**, pp. 2901–2906

Scientific Summary This paper extends the performance analysis of serial consensus to third- and higher-order systems. The notion of scalable performance is refined to account for amplification from external disturbances to transient deviations. It is shown that a broad class of non-identical, spatially distributed load disturbances can be rejected while maintaining uniformly bounded transients.

The third-order serial consensus system is especially relevant for vehicle platooning, where local integrator feedback is commonly used to compensate for disturbances such as slope and drag. The proposed controller takes the form:

$$u(t,x) = u_{\text{ref}}(t) - q_{\text{vel}} L\dot{x}(t) - q_{\text{pos}} L^2 x(t) - q_{\text{int}} L \int_0^t L^2 x(\tau)\, d\tau,$$

where $q_{\text{vel}}, q_{\text{pos}}, q_{\text{int}} > 0$ are design parameters. This proportional-integral structure preserves locality while enabling scalable disturbance rejection.

Numerical simulations confirm the effectiveness of the proposed protocol, demonstrating disturbance rejection without compromising transient performance.

Paper IV

> J. Hansson and E. Tegling (2025b). "Compositional design for time-varying and nonlinear coordination". *Preprint: arXiv:2504.07226*, in preparation for journal submission.

Scientific Summary This paper generalizes the serial consensus framework to accommodate nonlinear and time-varying dynamics, addressing practical challenges such as actuation delays, asynchronous measurements, and saturation constraints. The proposed control structure is given by:

$$\left(\frac{\mathrm{d}}{\mathrm{d}t} + \mathcal{L}_n\right) \circ \cdots \circ \left(\frac{\mathrm{d}}{\mathrm{d}t} + \mathcal{L}_1\right) x = 0,$$

where each $\mathcal{L}_k$ is a (possibly nonlinear or time-varying) Laplacian-like operator.

Similar to the linear case, this compositional consensus protocol guarantees asymptotic n^{th}-order consensus, provided that each first-order subsystem achieves consensus and mild regularity conditions are satisfied.

The framework supports a variety of use cases, including:

i) Consensus under elementwise saturation constraints;

ii) integration of time-delayed GPS data to enhance convergence;

iii) coordination with time-varying gain matrices.

Several examples are presented to illustrate the versatility of the proposed design. Together, these results show that high-order consensus with nonlinear and uncertain dynamics can be achieved within the proposed compositional framework.

Paper V

J. Hansson and E. Tegling (2022). "Input-output pseudospectral bounds for transient analysis of networked and high-order systems". In: *Proc. IEEE Conf. Decis. Control*, pp. 7497–7503

Scientific Summary This paper derives new bounds on the worst-case transient behavior of linear systems, expressed as:

$$\sup_{t\geq 0} \|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\|,$$

and relates this quantity to the frequency-domain behavior of the transfer matrix:

$$\|\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}\|.$$

The proof techniques are rooted in the theory of ε-pseudospectra, which are extended and adapted to an input-output setting.

These bounds are particularly useful for analyzing high-order networked systems, such as vehicle platooning models governed by:

$$s^2X(s)+sL_1X(s)+L_0X(s)=BU(s), \quad Y(s)=CX(s).$$

In contrast to classical spectral methods, the proposed pseudospectral bounds explicitly capture transient amplification and robustness limitations arising from double integrators and scaling effects in large-scale systems.

Contribution Statement

All included papers were authored by J. Hansson under the supervision and guidance of E. Tegling. J. Hansson was responsible for the main findings, mathematical analysis, and simulations. Both authors contributed to the writing, with the majority of the text drafted by J. Hansson. Final revisions and editing were carried out collaboratively.

4.2 Additional Publications

In addition to the five papers included in this thesis, the author has contributed to the following peer-reviewed publications:

J. Hansson, M. Svensson, A. Theorin, E. Tegling, K. Soltesz, T. Hägglund, and K. J. Åström (2021). "Next generation relay autotuners – analysis and implementation". In: *IEEE Conf. Control Technol. Appl.,* pp. 1075–1082

H. Pigot, J. Hansson, A. Paskevicius, Q. Liao, T. Sjöberg, S. Steen, and K. Soltesz (2021). "Identification of cardiac afterload dynamics from data". *IFAC-PapersOnLine* **54**:15, pp. 508–513

M. Lundh, A. Theorin, T. Hägglund, J. Hansson, M. Svensson, K. J. Åström, and K. Soltesz (2021). "Model optimization for autotuners in industrial control systems". In: *Proc. IEEE Int. Conf. Emerg. Technol. Fact. Autom.,* pp. 1–4

J. Hansson, A. Govaert, R. Pates, E. Tegling, and K. Soltesz (2022). "Limitations of time-delayed case isolation in heterogeneous sir models". In: *Proc. Am. Control Conf.,* pp. 2994–2999

F. Agner, J. Hansson, P. Kergus, A. Rantzer, S. Tarbouriech, and L. Zaccarian (2024). "Decentralized pi-control and anti-windup in resource sharing networks". *Eur. J. Control* **80**, p. 101049

The following extended abstract was also authored by J. Hansson:

J. Hansson and O. Kjellqvist (2024). "Magnitude-feedback control: a study on integrator stabilization". In: *Proc. Int. Symp. Math. Theory Netw. Syst.,* vol. 8, pp. 387–389

5

Discussion

This chapter discusses the implications and relevance of the proposed control frameworks, highlighting their strengths, practical applications, limitations, and potential directions for future research. It also summarizes the core findings and their significance within the broader fields of multi-agent coordination and scalable control.

5.1 Implication and Relevance to the Field

The proposed serial and compositional consensus frameworks offer several advantages that make them both theoretically robust and practically applicable. The key contributions include:

Simple and Easily Implementable Control Laws The serial and compositional designs yield structured yet simple control laws suitable for large-scale networked systems. Preliminary evaluations [Oorschot et al., 2025] indicate strong performance in realistic settings, including vehicular formations. These benefits come at the cost of modest additional signaling requirements, such as message passing or two-hop neighborhood measurements.

Robustness to Model Uncertainties The protocols offer provable robustness to parametric uncertainty, bounded disturbances, and model mismatches. Stability is guaranteed even under unmodeled dynamics—an essential feature for deployment in physical systems.

Adaptability to Practical Constraints The framework is highly flexible and can accommodate actuator saturations, time delays, asynchronous measurements, and heterogeneous agent dynamics. This adaptability makes the method well-suited to diverse real-world settings where ideal assumptions often fail.

Fuel Efficiency in Vehicle Platooning The ability to maintain robust formations independent of velocity is a major advantage in platooning applications. This feature enables aerodynamic drag reduction and potential fuel savings, improving energy efficiency without sacrificing stability or performance.

New Pseudospectra-Based Tools for Transient Analysis The thesis introduces novel analytical tools based on pseudospectra for assessing transient performance in high-order systems. These tools provide improved insights into system behavior, especially in the presence of structured uncertainties, and can be applied to a broader class of dynamical systems beyond those considered in this thesis.

5.2 Limitations

While the coordination frameworks developed in this thesis demonstrate strong scalability and robustness properties, several limitations remain.

Idealized System Model The framework assumes that each agent follows an ideal n^{th}-order integrator model. Although robustness results partially address model mismatches and unmodeled dynamics, this abstraction limits direct applicability to systems with more complex, heterogeneous, or uncertain behavior.

Implementation in very High-Order Systems In the general nonlinear or time-varying case, the number of nested control terms can grow rapidly with system order n, unless structural constraints—such as commuting operators—are imposed. This may limit practical deployment in very high-order systems.

Growth of Communication Neighborhood The locality of each agent's control law depends on the underlying graph structure. Without careful design (e.g., tree-like or sparse graphs), the required communication neighborhood can expand significantly with increasing n. In densely connected or small-world networks, this may result in nontrivial communication overhead.

Assumption of Full Cooperation The protocols rely on full participation and cooperation from all agents. In real-world applications such as vehicular formations, some agents may act passively, fail to share information, or be constrained by communication or safety policies. These cases are not explicitly handled in the current framework. Similarly, the thesis does not address practical challenges related to such cooperation, e.g. economic incentives, communication standards or safety and security concerns.

Lack of Stochastic Modeling The framework does not currently address noise, probabilistic agent behavior, or communication uncertainty. While deterministic robustness is established, stochastic performance and reliability remain open challenges.

Limited Treatment of Dynamic Topology Although the compositional framework supports time-varying dynamics, it does not rigorously address highly dynamic or adversarial switching in the communication graph. Existing results rely on assumptions such as persistent connectivity or smooth switching.

No Global Performance Optimization The protocols ensure local stability and robustness but are not optimized with respect to global performance criteria such as $\mathcal{H}_2$, $\mathcal{H}_\infty$, or $\mathcal{L}_\infty$ norms. The thesis does not propose a systematic approach to achieving global optimality within the compositional design paradigm.

5.3 Future Work

This thesis raises several new questions and directions for future research. Promising avenues include:

Observers in Place of Message Passing While the serial consensus framework currently relies on message passing to compute higher-order terms, an alternative approach based on local observers could be explored. Designing distributed observers that infer derivative or relative state information from local measurements may reduce communication requirements and improve robustness to communication losses.

Optimizing Coherence for Noise Rejection The class of stabilizing control laws introduced in this thesis creates new possibilities for optimizing network coherence in the presence of measurement noise. Investigating trade-offs between robustness, coherence, and scalability could yield improved designs for systems subject to stochastic disturbances.

Large-Scale Implementation and Testing A natural next step is to implement and evaluate the proposed control strategies in large-scale experiments. Applications such as autonomous vehicle platooning or robotic swarms offer opportunities to assess practical constraints, performance bottlenecks, and integration with sensing and communication infrastructure.

Optimization-Based Control Design The serial consensus structure provides a parameterized class of scalable control laws. Leveraging convex or nonlinear optimization to tune gains and feedback operators could improve performance in uncertain or time-varying environments while preserving locality and robustness.

Characterization of Scalable Performance This thesis establishes sufficient conditions for scalable performance, but a complete theoretical characterization remains open. Identifying necessary conditions—or establishing fundamental limits—would deepen understanding of when scalable performance is achievable in decentralized control systems.

5.4 Conclusions

This thesis has introduced novel design methods for coordinating networked systems. The proposed approach begins by specifying closed-loop dynamics with desirable properties and then systematically derives the corresponding control laws. By construction, the resulting protocols rely solely on local and relative feedback.

In the linear and time-invariant setting—captured by the serial consensus framework—stability, performance, and robustness guarantees have been explicitly established. In the nonlinear and time-varying setting—captured by the compositional consensus framework—stability guarantees have been proved under general conditions. Special emphasis has been placed on scalable performance with respect to transient behavior. The thesis develops scalable performance bounds, particularly in the context of string stability, and introduces new analytical tools—based on pseudospectra—to evaluate and predict worst-case transient amplification in non-normal dynamical systems.

While the foundational results presented here show strong potential, several directions remain open for future exploration. These include:

i) Incorporating observers to reduce or eliminate additional communication requirements;

ii) characterizing scalable performance for general nonlinear and time-varying consensus systems;

iii) implementing and validating the proposed methods in real-world networked systems.

Bibliography

Abolfazli, E., B. Besselink, and T. Charalambous (2023). "Minimum time headway in platooning systems under the mpf topology for different wireless communication scenario". *IEEE Trans. Intell. Transp. Syst.* **24**:4, pp. 4377–4390.

Agner, F., J. Hansson, P. Kergus, A. Rantzer, S. Tarbouriech, and L. Zaccarian (2024). "Decentralized pi-control and anti-windup in resource sharing networks". *Eur. J. Control* **80**, p. 101049.

Amirkhani, A. and A. H. Barshooi (2022). "Consensus in multi-agent systems: a review". *Artif. Intell. Rev.* **55**:5, pp. 3897–3935.

Asllani, M., R. Lambiotte, and T. Carletti (2018). "Structure and dynamical behavior of non-normal networks". *Science advances* **4**:12, eaau9403.

Baggio, G., V. Rutten, G. Hennequin, and S. Zampieri (2020). "Efficient communication over complex dynamical networks: the role of matrix non-normality". *Sci. Adv.* **6**:22, eaba2282.

Bamieh, B., M. R. Jovanovic, P. Mitra, and S. Patterson (2012). "Coherence in large-scale networks: dimension-dependent limitations of local feedback". *IEEE Trans. Autom. Control* **57**:9, pp. 2235–2249.

Barooah, P. and J. Hespanha (2005). "Error amplification and disturbance propagation in vehicle strings with decentralized linear control". In: *Proc. IEEE Conf. Decis. Control*, pp. 4964–4969.

Barooah, P., P. G. Mehta, and J. P. Hespanha (2009). "Mistuning-based control design to improve closed-loop stability margin of vehicular platoons". *IEEE Trans. Autom. Control* **54**:9, pp. 2100–2113.

Bernardo, C., C. Altafini, A. Proskurnikov, and F. Vasca (2024). "Bounded confidence opinion dynamics: a survey". *Automatica* **159**, p. 111302.

Besselink, B. and S. Knorn (2018). "Scalable input-to-state stability for performance analysis of large-scale networks". *IEEE Control Syst. Lett.* **2**:3, pp. 507–512.

Box, G. E. (1976). "Science and statistics". *J. Am. Stat. Assoc.* **71**:356, pp. 791–799.

Buser, P. (1982). "A note on the isoperimetric constant". In: *Ann. Sci. Éc. Norm. Supér.* Vol. 15. 2, pp. 213–230.

Chu, D., C. Zhao, R. Wang, Q. Xiao, W. Wang, and D. Cao (2024). "A survey of multi-vehicle consensus in uncertain networks for autonomous driving". *IEEE Trans. Intell. Transp. Syst.* **25**:12, pp. 19319–19341.

Chu, K.-C. (1974a). "Optimal dencentralized regulation for a string of coupled systems". *IEEE Trans. Autom. Control* **19**:3, pp. 243–246.

Chu, K.-c. (1974b). "Decentralized control of high-speed vehicular strings". *Transp. Sci.* **8**:4, pp. 361–384.

DeGroot, M. H. (1974). "Reaching a consensus". *J. Am. Stat. Assoc.* **69**:345, pp. 118–121.

Feintuch, A. and B. Francis (2012). "Infinite chains of kinematic points". *Automatica* **48**:5, pp. 901–908.

Feng, S., Y. Zhang, S. E. Li, Z. Cao, H. X. Liu, and L. Li (2019). "String stability for vehicular platoon control: definitions and analysis methods". *Annu. Rev. Control* **47**, pp. 81–97.

Friedkin, N. E. and E. C. Johnsen (1990). "Social influence and opinions". *J. Math. Sociol.* **15**:3-4, pp. 193–206.

Gunter, G., D. Gloudemans, R. E. Stern, S. McQuade, R. Bhadani, M. Bunting, M. L. Delle Monache, R. Lysecky, B. Seibold, J. Sprinkle, B. Piccoli, and D. B. Work (2021). "Are commercially implemented adaptive cruise control systems string stable?" *IEEE Trans. Intell. Transp. Syst.* **22**:11, pp. 6992–7003.

Hansson, J., A. Govaert, R. Pates, E. Tegling, and K. Soltesz (2022). "Limitations of time-delayed case isolation in heterogeneous sir models". In: *Proc. Am. Control Conf.,* pp. 2994–2999.

Hansson, J. and O. Kjellqvist (2024). "Magnitude-feedback control: a study on integrator stabilization". In: *Proc. Int. Symp. Math. Theory Netw. Syst.,* vol. 8, pp. 387–389.

Hansson, J., M. Svensson, A. Theorin, E. Tegling, K. Soltesz, T. Hägglund, and K. J. Åström (2021). "Next generation relay autotuners – analysis and implementation". In: *IEEE Conf. Control Technol. Appl.,* pp. 1075–1082.

Hansson, J. and E. Tegling (2022). "Input-output pseudospectral bounds for transient analysis of networked and high-order systems". In: *Proc. IEEE Conf. Decis. Control*, pp. 7497–7503.

Hansson, J. and E. Tegling (2023). "A closed-loop design for scalable high-order consensus". In: *Proc. IEEE Conf. Decis. Control*, pp. 7388–7394.

Hansson, J. and E. Tegling (2024). "Performance bounds for multi-vehicle networks with local integrators". *IEEE Control Syst. Lett.* **8**, pp. 2901–2906.

Hansson, J. and E. Tegling (2025a). "Closed-loop design for scalable performance of vehicular formations". *IEEE Trans. Control Netw. Syst.*, pp. 1–10.

Hansson, J. and E. Tegling (2025b). "Compositional design for time-varying and nonlinear coordination". *Preprint: arXiv:2504.07226*.

Herman, I., S. Knorn, and A. Ahlén (2017a). "Disturbance scaling in bidirectional vehicle platoons with different asymmetry in position and velocity coupling". *Automatica* **82**, pp. 13–20.

Herman, I., D. Martinec, Z. Hurák, and M. Sebek (2017b). "Scaling in bidirectional platoons with dynamic controllers and proportional asymmetry". *IEEE Trans. Autom. Control* **62**:4, pp. 2034–2040.

Huang, J., H. Fang, L. Dou, and J. Chen (2014). "An overview of distributed high-order multi-agent coordination". *IEEE/CAA J. Autom. Sin.* **1**:1, pp. 1–9.

Khalil, H. K. (2002). *Nonlinear Systems*. 3rd. Prentice Hall, Upper Saddle River, NJ.

Knorn, S., Z. Chen, and R. H. Middleton (2016). "Overview: collective control of multiagent systems". *IEEE Trans. Control Netw. Syst.* **3**:4, pp. 334–347.

Kuramoto, Y. (1975). "Self-entrainment of a population of coupled non-linear oscillators". In: *Proc. Int. Symp. Math. Probl. Theor. Phys.,* Springer, pp. 420–422.

Levine, W. and M. Athans (1966). "On the optimal error regulation of a string of moving vehicles". *IEEE Trans. Autom. Control* **11**:3, pp. 355–361.

Li, C.-J., G.-P. Liu, P. He, F. Deng, and J. Cao (2024). "Consensus of multiple high-order integrator agents with time-varying connectivity and delays: protocols using only relative states". *IEEE Trans. Syst., Man, Cybern.: Syst.* **54**:5, pp. 2663–2675.

Lin, Z., B. Francis, and M. Maggiore (2005). "Necessary and sufficient graphical conditions for formation control of unicycles". *IEEE Trans. Autom. Control* **50**:1, pp. 121–127.

Lu, M. and L. Liu (2017). "Distributed feedforward approach to cooperative output regulation subject to communication delays and switching networks". *IEEE Trans. Autom. Control* **62**:4, pp. 1999–2005.

Lundh, M., A. Theorin, T. Hägglund, J. Hansson, M. Svensson, K. J. Åström, and K. Soltesz (2021). "Model optimization for autotuners in industrial control systems". In: *Proc. IEEE Int. Conf. Emerg. Technol. Fact. Autom.,* pp. 1–4.

Lyu, J., J. Qin, D. Gao, and Q. Liu (2016). "Consensus for constrained multi-agent systems with input saturation". *Int. J. Robust Nonlinear Control* **26**:14, pp. 2977–2993.

Melzer, S. and B. Kuo (1971). "Optimal regulation of systems described by a countably infinite number of objects". *Automatica* **7**:3, pp. 359–366.

Middleton, R. H. and J. H. Braslavsky (2010). "String instability in classes of linear time invariant formation control with limited communication range". *IEEE Trans. Autom. Control* **55**:7, pp. 1519–1530.

Moreau, L. (2005). "Stability of multiagent systems with time-dependent communication links". *IEEE Trans. Autom. Control* **50**:2, pp. 169–182.

Moreau, L. (2004). "Stability of continuous-time distributed consensus algorithms". In: *Proc. IEEE Conf. Decis. Control*. Vol. 4, pp. 3998–4003.

Olfati-Saber, R. and R. M. Murray (2004). "Consensus problems in networks of agents with switching topology and time-delays". *IEEE Trans. Autom. Control* **49**:9, pp. 1520–1533.

Oorschot, T. van, M. Jeeninga, and E. Tegling (2025). *Experimental verification of a scalable protocol for vehicle platooning*. To appear in: *Proc. Eur. Control Conf.*

Pates, R., C. Lidström, and A. Rantzer (2017). "Control using local distance measurements cannot prevent incoherence in platoons". In: *Proc. IEEE Conf. Decis. Control*, pp. 3461–3466.

Patterson, S. and B. Bamieh (2010). "Leader selection for optimal network coherence". In: *Proc. IEEE Conf. Decis. Control*, pp. 2692–2697.

Pigot, H., J. Hansson, A. Paskevicius, Q. Liao, T. Sjöberg, S. Steen, and K. Soltesz (2021). "Identification of cardiac afterload dynamics from data". *IFAC-PapersOnLine* **54**:15, pp. 508–513.

Radmanesh, A., A. Naghash, and A. Mohamadifard (2017). "Optimal distributed control of multi agents: generalization of consensus algorithms for high-order state derivatives of siso and mimo systems". In: *Proc. Int. Conf. Control Autom. Robot.,* pp. 606–611.

Ren, W., K. L. Moore, and Y. Chen (2006). "High-order consensus algorithms in cooperative vehicle systems". In: *Proc. IEEE Int. Conf. Netw. Sens. Control*, pp. 457–462.

Ren, W., K. L. Moore, and Y. Chen (2007). "High-order and model reference consensus algorithms in cooperative control of multi-vehicle systems". *J. Dyn. Syst. Meas. Control* **129**:5, pp. 678–688.

Ren, W., R. W. Beard, and T. W. McLain (2005). "Coordination variables and consensus building in multiple vehicle systems". In: Kumar, V. et al. (Eds.). *Cooperative Control: A Post-Workshop Volume 2003 Block Island Workshop on Cooperative Control*. Springer, Berlin, pp. 171–188.

Rezaee, H. and F. Abdollahi (2015). "Average consensus over high-order multiagent systems". *IEEE Trans. Autom. Control* **60**:11, pp. 3047–3052.

Seiler, P., A. Pant, and K. Hedrick (2004). "Disturbance propagation in vehicle strings". *IEEE Trans. Autom. Control* **49**:10, pp. 1835–1842.

Silva, G. F., A. Donaire, R. Middleton, A. McFadyen, and J. Ford (2025). "Scalable input-to-state stability of nonlinear interconnected systems". *IEEE Trans. Autom. Control* **70**:3, pp. 1824–1834.

Stern, R. E., S. Cui, M. L. Delle Monache, R. Bhadani, M. Bunting, M. Churchill, N. Hamilton, R. Haulcy, H. Pohlmann, F. Wu, B. Piccoli, B. Seibold, J. Sprinkle, and D. B. Work (2018). "Dissipation of stop-and-go waves via control of autonomous

vehicles: field experiments". *Transp. Res. Part C Emerg. Technol.* **89**, pp. 205–221.

Strogatz, S. H. (2000). "From kuramoto to crawford: exploring the onset of synchronization in populations of coupled oscillators". *Phys. D: Nonlinear Phenom.* **143**:1, pp. 1–20.

Stüdli, S., M. M. Seron, and R. H. Middleton (2017a). "From vehicular platoons to general networked systems: string stability and related concepts". *Annu. Rev. Control* **44**, pp. 157–172.

Stüdli, S., M. M. Seron, and R. H. Middleton (2017b). "Vehicular platoons in cyclic interconnections with constant inter-vehicle spacing". *IFAC-PapersOnLine* **50**:1, pp. 2511–2516.

Swaroop, D. and J. Hedrick (1996). "String stability of interconnected systems". *IEEE Trans. Autom. Control* **41**:3, pp. 349–357.

Tegling, E., B. Bamieh, and H. Sandberg (2023). "Scale fragilities in localized consensus dynamics". *Automatica* **153**, p. 111046.

Tegling, E. and H. Sandberg (2017). "On the coherence of large-scale networks with distributed pi and pd control". *IEEE Control Syst. Lett.* **1**:1, pp. 170–175.

Tian, Y.-P. and Y. Zhang (2012). "High-order consensus of heterogeneous multi-agent systems with unknown communication delays". *Automatica* **48**:6, pp. 1205–1212.

Trefethen, L. and M. Embree (2005). *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*. Princeton University Press.

Trindade, P., P. Batista, and R. Cunha (2025). "Delay-margin design in consensus for high-order integrator agents". *Automatica* **174**, p. 112173.

Trindade, P., R. Cunha, and P. Batista (2024). "Optimal consensus for high-order integrator agents". In: *Proc. IEEE Conf. Decis. Control*, pp. 6957–6962.

Ureña, R., G. Kou, Y. Dong, F. Chiclana, and E. Herrera-Viedma (2019). "A review on trust propagation and opinion dynamics in social networks and group decision making frameworks". *Inf. Sci.* **478**, pp. 461–475.

Vegamoor, V. K., S. Darbha, and K. R. Rajagopal (2019). "A review of automatic vehicle following systems". *J. Indian Inst. Sci.* **99**, pp. 567–587.

Yu, W., G. Chen, W. Ren, J. Kurths, and W. X. Zheng (2011). "Distributed higher order consensus protocols in multiagent dynamical systems". *IEEE Trans. Circuits Syst. I: Regul. Pap.* **58**:8, pp. 1924–1932.

Paper I

A Closed-Loop Design for Scalable High-Order Consensus

Jonas Hansson **Emma Tegling**

Abstract

This paper studies the problem of coordinating a group of n^{th}-order integrator systems. As in the well-studied conventional consensus problem, we consider linear and distributed control with only local and relative measurements. We propose a closed-loop dynamic that we call *serial consensus* and prove it achieves n^{th}-order consensus regardless of model order and underlying network graph. This alleviates an important scalability limitation in conventional consensus dynamics of order $n \geq 2$, whereby they may lose stability if the underlying network grows. The distributed control law which achieves the desired closed loop dynamics is shown to be localized and obey the limitation to relative state measurements. Furthermore, through use of the small-gain theorem, the serial consensus system is shown to be robust to both model and feedback uncertainties. We illustrate the theoretical results through examples.

1. Introduction

Properties of dynamical systems over networks have been a subject of significant research over the last two decades. A problem of interest is the coordination of agents in a network through localized feedback, leading to the prototypical distributed consensus dynamics studied early on by [Fax and Murray, 2004; Olfati-Saber and Murray, 2004; Jadbabaie et al., 2003]. Over the years, it has become clear that the structural constraints imposed by the network topology in consensus problems often lead to fundamentally poor dynamic behaviors in large networks. This concerns controllability [Pasqualetti et al., 2014], performance [Bamieh et al., 2012; Siami and Motee, 2016] and disturbance propagation [Swaroop and Hedrick, 1996; Seiler et al., 2004], but, as recently highlighted in [Tegling et al., 2023], also stability. The poor stability properties characterized in earlier work [Tegling et al., 2023] (which motivate the present work) apply to higher-order consensus. Here the local dynamics of each agent is modeled as an n^{th}-order integrator, with $n \geq 2$, and the control is a weighted average of neighbors' relative states. This is a theoretical generalization of first-order consensus [Jiang et al., 2009], but is also relevant in practice. For example, a model where $n = 3$ and thus has consensus in position, velocity and acceleration, can capture flocking behaviors [Ren et al., 2006].

More specifically, [Tegling et al., 2023] shows that conventional high-order consensus ($n \geq 3$) is not *scalably stable* for many growing graph structures. When the network grows beyond a certain size, stability is lost. The same holds for second-order consensus ($n = 2$) in, for example, directed ring graphs, as also observed in [Stüdli et al., 2017]. To address this lack of scalable stability we propose an alternative generalization of the first-order consensus dynamics, which achieves consensus regardless of underlying network graph and model order n, thereby also ensuring scalable stability.

To illustrate our proposed controller, consider first the conventional second-order consensus system. Here the controller $u(t) = -L_1\dot{x}(t) - L_0x(t) + u_{\text{ref}}(t)$, with $L_{0,1}$ being weighted graph Laplacians, is used to achieve the closed loop

$$\ddot{x}(t) = -L_1\dot{x}(t) - L_0x + u_{\text{ref}}(t). \tag{1}$$

While for first-order consensus ($\dot{x} = Lx + u_{\text{ref}}$), a sufficient condition for convergence to consensus is that the graph underlying the graph Laplacian L contains a connected spanning tree [Ren et al., 2007]. However, this no longer suffices when $n \geq 2$ as in (1). Therefore, we instead propose the following controller $u(t) = -(L_2 + L_1)\dot{x}(t) - L_2L_1x(t) + u_{\text{ref}}(t)$. The reason for this choice of controller is best illustrated by considering the resulting closed loop in the Laplace domain:

$$(sI + L_2)(sI + L_1)X(s) = U_{\text{ref}}. \tag{2}$$

For this system, like for the first-order case, it is sufficient that the graphs underlying L_1 and L_2 each contain a connected spanning tree for the system to eventually

coordinate in both x and its derivative $\dot{x}$ (regardless of network size!). This closed loop system, which we will call *serial consensus*, thus mimics one core property of the standard consensus protocol, and can also be generalized to any order n.

The main results of this paper elucidate key properties of the proposed n^{th}-order serial consensus. The controller is proven to remain localized (within an n-hop neighborhood) and implementable through relative measurements. We also prove that the closed loop will achieve consensus in all n states. Furthermore, we study the robustness of the proposed closed loop and show that the system will still coordinate when subject to unstructured uncertainty, whose permissible size is independent of the network size. The beneficial properties of the form (2) (generalized to any order n) are thus not contingent on an idealized implementation.

The remainder of this paper is organized as follows. We begin by introducing the n^{th}-order consensus model and defining our choice of control structure. Subsequently, we define and motivate the serial consensus system. In Section 3 and 4, we provide proofs for the stability and robustness of the serial consensus system, respectively. The main results are then illustrated through examples in Section 5. Finally, Section 6 offers our conclusions.

2. Problem Setup

We start by introducing some notation and graph theory before introducing the general n^{th}-order consensus problem for which we propose the new serial consensus setup. We discuss its properties and end with some useful definitions.

2.1 Network Model and Definitions

Let $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ denote a graph of size $N = |\mathcal{V}|$. The set $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ denotes the set of edges. The graph can be equivalently represented by the adjacency matrix $W \in \mathbb{R}^{N \times N}$ where $w_{i,j} > 0 \iff (j,i) \in \mathcal{E}$. The graph is called *undirected* if $W = W^T$. The graph contains a *connected spanning tree* if for some $i \in \mathcal{V}$ there is a path from i to any other vertex $j \in \mathcal{V}$.

Associated with a weighted graph we have the weighted graph Laplacian $L(\mathcal{G})$, defined as

$$[L(\mathcal{G})]_{i,j} = \begin{cases} -w_{i,j}, & \text{if } i \neq j \\ \sum_{k \neq i} w_{i,k}, & \text{if } i = j \end{cases}.$$

The graph dependence is omitted when clear from context. Under the condition that that the graph generating the graph Laplacian contains a connected spanning tree, L will have a simple and unique eigenvalue at 0 and the remaining eigenvalues will lie strictly in the right half plane (RHP).

We will also consider networks with a growing number of nodes. These are then described by a graph family $\{\mathcal{G}_N\}_{N \to \infty}$, where N is the size of the growing network.

The space of all proper, real rational, and stable transfer matrices will be denoted $\mathcal{RH}_\infty$. We will use $\|\cdot\|_{\mathcal{H}_\infty}$ for the $\mathcal{H}_\infty$ norm, following the notation in [Zhou and Doyle, 1998].

2.2 n^{th}-Order Consensus

Let the system be modeled as N agents with identical n^{th}-order integrator dynamics, i.e.

$$\frac{\mathrm{d}^n x_i(t)}{\mathrm{d}t^n} = u_i(t), \tag{3}$$

for all $i \in \mathcal{V}$. We will use the convention $x_i^{(0)}(t) = x_i(t)$ and $x_i^{(k)}(t) = \frac{\mathrm{d}^k}{\mathrm{d}t^k} x_i(t)$ to denote time derivatives. When clear, we may omit the time argument for brevity.

In this paper, we consider the problem of synchronizing agents to achieve a state of consensus, formally defined as:

Definition 1—n^{th}-order consensus

The multi-agent system (3) achieves n^{th}-order consensus if $\lim_{t\to\infty} |x_i^{(k)}(t) - x_j^{(k)}(t)| = 0$*, for all $i, j \in \mathcal{V}$ and $k \in \{0, 1 \ldots, n-1\}$.*

2.3 Control Structure

A linear state feedback controller of (3) can be written as

$$u(t) = u_{\text{ref}}(t) - \sum_{k=0}^{n-1} A_k x^{(k)}(t), \tag{4}$$

where $u_{\text{ref}}(t) \in \mathbb{R}^N$ is a feedforward term and $A_k \in \mathbb{R}^{N\times N}$ represents the feedback of the k^{th} derivative. We will restrict this class of controllers in three ways. The controllers

i) can only use *relative* feedback;

ii) have a limited gain;

iii) depend on the local neighborhood of each agent.

The constraint for relative feedback can be expressed as $A_k \mathbf{1}_N = 0$ for all k. Meanwhile, a limited gain can be represented by requiring that $\|A_k\|_\infty \leq c$. To capture the notion of locality, consider the adjacency matrix W representing the communication and measurement structure, which we here assume to be the same. That is, if $W_{i,j} = 1$, then agent i can directly receive or measure the relative distance to agent j. Next, consider the non-negative matrix W^q. This matrix has the property that $[W^q]_{i,j} \neq 0$ if and only if there is a path of length q from agent j to agent i. Thus, if we want the controller to only depend on information that is at most q steps away from each agent the following implication should hold: $\sum_{k=0}^{q} W_{i,j}^k = 0 \implies [A_k]_{i,j} = 0$. Putting all the conditions together gives us a family of controllers that we will consider in this paper:

Definition 2—q-step implementable relative feedback
A relative feedback controller of the form (4) is q-step implementable with respect to the adjacency matrix W and gain $c > 0$ if $A_k \in \mathcal{A}^q(W,c)$ for all k, where

$$\mathcal{A}^q(W,c) = \left\{ A \;\middle|\; \begin{matrix} \left[\Sigma_{k=0}^{q} W^k\right]_{i,j} = 0 \implies A_{i,j} = 0, \\ A\mathbf{1}_N = 0,\ \|A\|_\infty \le c \end{matrix} \right\}.$$

The conventional controller for achieving n^{th}-order consensus can be realized as (4) where each A_k is given by a graph Laplacian, e.g., $A_k = L_k \in \mathcal{A}^1(W,c)$.

2.4 A Novel Design: Serial Consensus

We propose the following controller of (3), expressed in the Laplace domain, to achieve n^{th}-order consensus

$$U(s) = U_{\text{ref}}(s) + \left(s^n I - \prod_{k=1}^{n} (sI + L_k) \right) X(s), \tag{5}$$

where L_k are graph Laplacians and U_{ref} is the transformed reference signal. In this case, it is more instructive to consider the closed-loop dynamics, which take the following form:

Definition 3—n^{TH}-order serial consensus system
For all $k \in \{1,2\ldots,n\}$, let L_k be a weighted and directed graph Laplacian. The n^{th}-order serial consensus system is then

$$\left(\prod_{k=1}^{n} (sI + L_k) \right) X(s) = U_{\text{ref}}(s). \tag{6}$$

We refer to this form as serial consensus because the same closed-loop dynamics can be realized through interconnecting n first-order consensus systems in series.

The closed-loop dynamics in (6) can also be transformed to state-space form by introducing the alternative variables Ξ_k with the corresponding states ξ_k. These relate to X through $\Xi_1 = X(s)$, $\Xi_k = (sI + L_{k-1})\Xi_{k-1}$ for $k \in \{2,\ldots,n-1\}$, and $s\Xi_n = -L_n\Xi_n + U_{\text{ref}}$. This leads to the following continuous-time state-space representation

$$\begin{bmatrix} \dot{\xi}_1 \\ \dot{\xi}_2 \\ \vdots \\ \dot{\xi}_{n-1} \\ \dot{\xi}_n \end{bmatrix} = \underbrace{\begin{bmatrix} -L_1 & I & & \\ & -L_2 & \ddots & \\ & & \ddots & I \\ & & & -L_n \end{bmatrix}}_{A} \begin{bmatrix} \xi_1 \\ \xi_2 \\ \vdots \\ \xi_{n-1} \\ \xi_n \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ u_{\text{ref}} \end{bmatrix}. \tag{7}$$

The serial consensus form has several advantages, which will be the focus of this paper. First, however, we show that it satisfies the controller constraints that we impose, as given by Definition 2. In other words, we will discuss how the closed-loop structure in (6) can be implemented on a network.

When analyzing the serial consensus controller of (5) we will utilize the following assumption on the graph structure.

Assumption 1
(Connected spanning tree) All graphs underlying the graph Laplacians L_k contain a connected spanning tree.

2.5 Implementing Serial Consensus

The following proposition ensures that the serial consensus system can be achieved by controlling the n^{th}-order integrator system (3) with an n-step implementable relative feedback controller as defined in Definition 2.

Proposition 1
Consider the n^{th}-order serial consensus as defined in (6). If each $L_k \in \mathcal{A}^1(W,c)$ for some constant c and adjacency matrix W, then the controller in (5) is an n-step implementable relative feedback controller with respect to W and a finite gain c'.

To prove this proposition, we first require the following two lemmas, whose proofs can be found in [Hansson and Tegling, 2023].

Lemma 1
If $A_1 \in \mathcal{A}^{q_1}(W,c_1)$ and $A_2 \in \mathcal{A}^{q_2}\ (W,c_2)$ then the sum satisfies $(A_1+A_2) \in \mathcal{A}^{\max(q_1,q_2)}\ (W,c_1+c_2)$.

Lemma 2
Let $A_1 \in \mathcal{A}^{q_1}(W,c_1)$ and $A_2 \in \mathcal{A}^{q_2}(W,c_2)$ then the product satisfies $(A_1A_2) \in \mathcal{A}^{q_1+q_2}(W,c_1c_2)$.

Now we can prove *Proposition* 1.

Proof. The serial consensus controller can be expanded to the matrix polynomial

$$\begin{aligned} U(s) &= U_{\text{ref}}(s) + \left(s^n I - \prod_{k=1}^{n} (sI + L_k) \right) X(s) \\ &= U_{\text{ref}}(s) + \left((s^n - s^n) I - \sum_{k=0}^{n-1} s^k A_k \right) X(s), \end{aligned}$$

for some matrices A_k. To show the proposition, we need to show that $A_k \in \mathcal{A}^q(W,c')$ for all $k = 0,\dots,n-1$, with $q \leq n$ and $c' < \infty$. Let

$$\mathcal{I}_k \quad = \quad \left\{ \alpha \;\middle|\; |\alpha| = n-k,\ \alpha \subset \{1,2,\dots,n\},\ i<j \implies \alpha(i) < \alpha(j) \right\}$$

denote all the ordered subsets of the range $[1,n]$ with size $n-k$. Then

$$A_k = \sum_{\alpha \in \mathcal{I}_k} \prod_{j \in \alpha} L_j, \text{ for all } k \in [0, n-1].$$

Since all $\alpha \in \mathcal{I}_k$ has $n-k$ elements we can show that $\prod_{j\in\alpha} L_j = B_\alpha \in \mathcal{A}^{n-k}(W, c^{n-k})$ by applying Lemma 2 recursively. Now we have a sum

$$A_k = \sum_{\alpha \in \mathcal{I}_k} B_\alpha.$$

The number of ordered subsets of the range $[1,n]$ with size $n-k$ is given by the binomial coefficients and therefore the size of $|\mathcal{I}_k| = \binom{n}{n-k}$. Applying Lemma 1 recursively shows that $A_k \in \mathcal{A}^{n-k}(W, \binom{n}{n-k} c^{n-k})$. Clearly, we have that $n-k \leq n$ and $\binom{n}{n-k} c^{n-k} \leq \binom{n}{\lceil n/2 \rceil} \max(c, c^n) < \infty$ for all k. Let $c' = \binom{n}{\lceil n/2 \rceil} \max(c, c^n)$ and then $A_k \in \mathcal{A}^n(W, c')$ holds true for all k. □

Example 1
For clarity, let us consider the controller for the case $n = 3$. Then the controller is

$$\begin{aligned} U(s) &= U_{\text{ref}}(s) + \left(s^3 I - \prod_{k=1}^{3} (sI + L_k) \right) X(s) \\ &= U_{\text{ref}}(s) - \left(s^2 (L_3 + L_2 + L_1) + s(L_3 L_2 + L_3 L_1 + L_2 L_1) + L_3 L_2 L_1 \right) X(s). \end{aligned}$$

Here, $A_0 = L_3 L_2 L_1$, $A_1 = L_3 L_2 + L_3 L_1 + L_2 L_1$, and $A_2 = L_3 + L_2 + L_1$. The proposition asserts that if L_1, L_2, and L_3 share a sparsity pattern and have bounded gains, then the resulting controller gains A_0, A_1, and A_2 will be sparse and have bounded gains.

Remark 1
Proposition 1 may be conservative. For instance, if W represents the complete graph, then any relative feedback controller would trivially be 1-step implementable.

3. Stability of Serial Consensus

In this section, we prove the stability of the serial consensus. Using this result, we further demonstrate that the serial consensus satisfies a notion of scalable stability.

3.1 Stability

Theorem 1
Consider the n^{th}-order serial consensus system as defined in Definition 3 under Assumption 1 and with $U_{\text{ref}} \in \mathcal{RH}_\infty$. The closed-loop dynamics have the following properties:

i) The poles of (6) are given by the union of the eigenvalues of $-L_k$.

ii) The solution achieves n^{th}-order consensus.

Proof. i Any square matrix can be unitarily transformed to upper-triangular form by the Schur traingularization theorem. Let $U_k L_k U_k^H = T_k$ be upper triangular. Then the block diagonal matrix $U = \mathrm{diag}(U_1, U_2, \dots U_n)$ is a unitary matrix that upper triangularizes A in (7). For any triangular matrix the eigenvalues lie on the diagonal and these will be the eigenvalues of each $-L_k$.

(ii) First, consider the closed loop dynamics of (6) which will be

$$X(s) = \left(\prod_{k=n}^{1} (sI + L_k)^{-1} \right) U_{\mathrm{ref}}(s).$$

Since U_{ref} is stable, we know that the limit $\lim_{s\to 0} U_{\mathrm{ref}}(s) = U_{\mathrm{ref}}(0)$ exists. To prove that the system achieves n^{th}-order consensus we want to show that

$$\lim_{t\to\infty} y(t) = \lim_{s\to 0} C(s)X(s) = 0$$

for some transfer matrix $C(s)$, which encodes the consensus states. But since the reference dependence is only related to $U_{\mathrm{ref}}(0)$, we can simplify the problem to only consider impulse responses. But the impulse response has the same transfer function as the initial value response where $\xi_n(0) = U_{\mathrm{ref}}(0)$. Therefore, WLOG, assume that $U_{\mathrm{ref}}(s) = 0$ and an arbitrary initial condition

$$\xi(0) = [\xi_1^T(0), \xi_2^T(0), \dots, \xi_n^T(0)]^T.$$

The solution of (7) is given by $\exp(At)\xi(0) = S\exp(J(A)t)S^{-1}\xi(0)$ where $J(A)$ is the Jordan normal form of A and S is an invertible matrix. From i and the diagonal dominance of the graph Laplacians we know that all eigenvalues of A lie in the left half plane. By Assumption 1 it follows that the zero eigenvalue for each L_k is simple. Now we prove that these n zero eigenvalues form a Jordan block of size n. Let e_k denote the k^{th} 1-block vector, e.g. $\mathbf{e}_1 = [\mathbf{1}_N^T, 0_N, \dots, 0_N]^T$ and $e_2 = [0_N, \mathbf{1}_N^T, \dots, 0_N]^T$. Then e_1 is an eigenvector, since $A\mathbf{e}_1 = 0$. For $k \in \{2, 3 \dots, n\}$ we have $Ae_k = e_{k-1}$ which implies that $A^k e_k = 0$. This shows that there is a Jordan block of size n with an invariant subspace spanned by the vectors e_k. Since all other eigenvectors make up an asymptotically stable invariant subspace, it follows that $\xi(t)$ will converge towards a solution in $\mathrm{span}(e_1, e_2 \dots, e_n)$ and thus $\lim_{t\to\infty} \xi_k(t) = \alpha_k(t)\mathbf{1}_N$. Now, since $x(t) = \xi_1(t)$, it follows that $\lim_{t\to\infty} x(t) = \alpha_1(t)\mathbf{1}_N$, and furthermore, since

$$\dot{\xi}_k = -L_k\xi_k + \xi_{k+1} \to \xi_{k+1} \text{ as } t \to \infty,$$

for $k \in \{1, \dots, n-1\}$, it follows that $\lim_{t\to\infty} x^{(k)}(t) = \alpha_{k+1}(t)1_N$, proving convergence to n^{th}-order consensus. □

The proposition shows that the stability analysis of the n^{th}-order serial consensus can be reduced to verifying that the n first-order consensus systems $\dot{x} = -L_k x$ achieve consensus. This is equivalent to determining whether the graphs underlying each L_k, all contain a connected spanning tree. Consequently, in conjunction with *Proposition* 1, this result demonstrates that n^{th}-order consensus is achievable using a local relative feedback controller with finite gain. Notably, this achievement is independent of the number of agents, thus ensuring scalability, which we will discuss next.

3.2 Scalable Stability

Coordinating a multi-agent system is inherently a decentralized problem where the goal for each agent is to coordinate with its nearest neighbors. However, when the controllers only depend on local measurements there is a possibility that controllers that manage to coordinate N agents stop doing so as the number of agents increases. More specifically, consider the growing graph family $\{\mathcal{G}_N\}$ and corresponding graph Laplacians $L(\mathcal{G}_N)$. Then we can consider the following notion of stability.

Definition 4—Scalable stability [Tegling et al., 2023, Def. 2.1]
A consensus control design is scalably stable if the resulting closed-loop system achieves consensus over any graph in the family $\{\mathcal{G}_N\}$.

In [Tegling et al., 2023] it was shown that for the 3$^{\text{rd}}$ and higher-order consensus problem with controller $A_k = a_k L(\mathcal{G}_N)$ in (4), the closed loop system will become unstable if the algebraic connectivity $\lambda_2(L(\mathcal{G}_N)) \to 0$ as $N \to \infty$. The serial consensus alleviates this scalability issue. Theorem 1 shows that the serial consensus will be stable, regardless of the network size, as long as the underlying graphs are sufficiently connected. The result is summarized in the following corollary.

Corollary 1
For any n, the controller (5) is scalably stable over any graph family $\{\mathcal{G}_N\}$ that underlies $L_k(\mathcal{G}_N)$, provided each $\mathcal{G}_N$ satisfies Assumption 1.

Remark 2
Note that, by Theorem 1, scalable stability is also achieved when the graph families underlying each L_k are different. This can even be achieved with $||L_k||_\infty$ being arbitrarily small.

4. Robustness of Serial Consensus

The controller proposed in (5) is a relative state-feedback controller, specifically designed to ensure that the closed loop system achieves n^{th}-order consensus as guaranteed through Theorem 1. However, the n^{th}-order integrator system may be an idealization of the system. Implementing the relative state feedback may require

observers to be fully realized, and there can be unmodeled dynamics. These potential sources of errors call for a robust controller. We will now present two theorems, which prove that the serial consensus is robust towards two different types of uncertainties.

4.1 Additive Perturbation

The following theorem asserts that the n^{th}-order serial consensus controller can handle additive perturbations.

THEOREM 2
Consider the n^{th}-order serial consensus system as defined in Definition 3, under Assumption 1, with $L_k = L$ for all k, and $L = L^T$. Then the perturbed system

$$(sI+L)^n X = U_{\text{ref}} + \left(\sum_{k=0}^{n} \Delta_k s^k L^{n-k}\right) X,$$

where $U_{\text{ref}}, \Delta_k \in \mathcal{RH}_\infty$, achieves n^{th}-order consensus if

$$\|\Delta_0\|_{\mathcal{H}_\infty} + \|\Delta_n\|_{\mathcal{H}_\infty} + \sum_{k=1}^{n-1} \|\Delta_k\|_{\mathcal{H}_\infty} \sqrt{\frac{k^k}{n^n}(n-k)^{n-k}} < 1.$$

Proof. First, note that the closed-loop system can be represented by the block diagram in Figure 1, which in turn can be simplified to Figure 2. Since U_{ref} is stable we can apply the small-gain theorem which asserts that $U(s)$ (as defined in the figures) will be stable if

$$\left\| \sum_{k=0}^{n} \Delta_k s^k L^{n-k} (sI+L)^{-n} \right\|_{\mathcal{H}_\infty} < 1.$$

Applying the triangle inequality and submultiplicativity on the left-hand side (LH) yields

$$LH \leq \sum_{k=0}^{n} \|\Delta_k\|_{\mathcal{H}_\infty} \left\| s^k L^{n-k} (sI+L)^{-n} \right\|_{\mathcal{H}_\infty}. \tag{8}$$

Since L is symmetric, it is possible to unitarily diagonalize it. Let $U = U^H$ denote one such unitary matrix. Then $L = U\Lambda U^H$ where Λ is a non-negative real diagonal matrix.

$$\|s^k L^{n-k}(sI+L)^{-n}\|_{\mathcal{H}_\infty} = \|s^k \Lambda^{n-k}(sI+\Lambda)^{-n}\|_{\mathcal{H}_\infty}.$$

For a diagonal matrix the singular values are given by the absolute value of the diagonal. Let, $\lambda > 0$ be an arbitrary positive constant. The maximum gain for each diagonal can then be calculated through

$$\max_{\omega} \left| \frac{\omega^k \lambda^{n-k}}{(j\omega+\lambda)^n} \right| = \sqrt{\max_{\omega} \frac{\omega^{2k}\lambda^{2n-2k}}{(\omega^2+\lambda^2)^n}}.$$

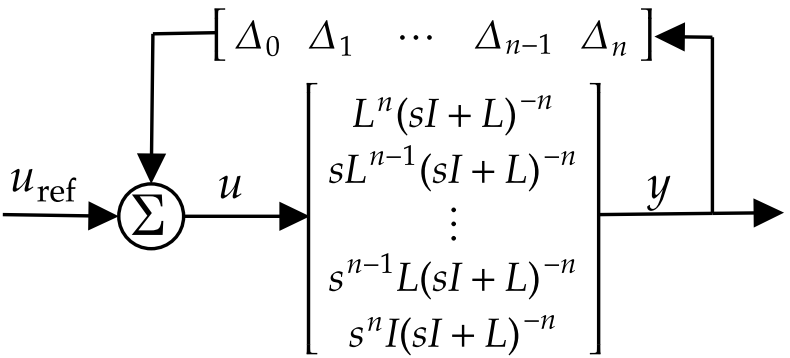

Figure 1. Block diagram illustrating the perturbation model in proof of Theorem 2.

The latter optimization problem is given by a continuous function and thus the derivative must be 0 at the maximum. Simple calculus shows that the optimum is found at $\omega^2 = \lambda^2 k/(n-k)$ for $k = 0, 1, \ldots n-1$ and at $\omega = \infty$ for $k = n$. Inserting yields

$$\max_{\omega} \left| \frac{\omega^k \lambda^{n-k}}{(j\omega + \lambda)^n} \right| = \begin{cases} \sqrt{\frac{k^k}{n^n}(n-k)^{n-k}} & \text{if } 0 < k < n \\ 1 & \text{else} \end{cases}.$$

In the case where $\lambda = 0$. Then we have for $k = 0, \ldots, n-1$

$$\max_{\omega} \left| \frac{\omega^k 0^{n-k}}{(j\omega + 0)^n} \right| = 0,$$

and for $k = n$

$$\max_{\omega} \left| \frac{\omega^n}{(j\omega + 0)^n} \right| = 1.$$

This is less restrictive than for $\lambda > 0$ and thus we can use the result for $\lambda > 0$. Plugging this into the upper bound of the *LH* (8) results in the sought inequality.

Finally, we must ensure that stability of the closed loop in Figure 2 implies n^{th}-order consensus. Since the transfer matrix from u to y in Figure 1 is stable it follows that $Y(s)$ will be stable. This means that we have shown the following $\lim_{t\to\infty} L^{n-k} x^{(k)}(t) = 0$. By Assumption 1 the 0 eigenvalue of L is unique and therefore 0 is a unique eigenvalue of L^{n-k} too. Subsequently, $\lim_{t\to\infty} x^{(k)}(t) \in \ker(L^{n-k})$. Since $L^{n-k}\mathbf{1}_N = 0$ it follows that $\lim_{t\to\infty} x^{(k)}(t) \in \text{span}(\mathbf{1}_N)$ and that the agents will reach consensus in all the $n-1$ first time derivatives and thus achieve n^{th}-order consensus. □

It is worth noting that the norm bound on the uncertainty blocks Δ_k is independent of the number of agents in the system. Therefore, the serial consensus implementation can be considered *scalably robust* in the sense that it allows equally sized perturbations, regardless of network size. This is not the case for localized conventional consensus, following the results in [Tegling et al., 2023].

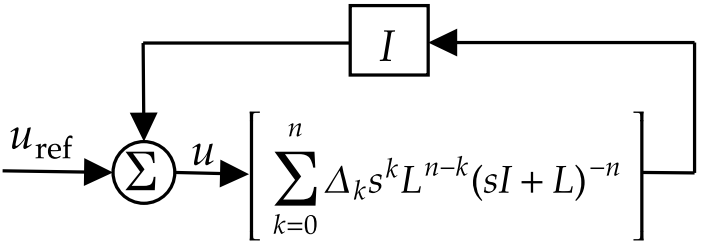

Figure 2. Block diagram illustrating the perturbation model in proof of Theorem 2.

4.2 Multiplicative Perturbation

It is also possible to see the closed-loop serial consensus system as a series of interconnected first-order systems. Therefore it is also interesting to consider the robustness with respect to the each factor. The following theorem gives a sufficient condition for the closed-loop system to achieve n^{th}-order consensus.

Theorem 3
Consider the n^{th}-order serial consensus system as defined in Definition 3, under Assumption 1, with $L_k = L_k^T$ for all k. Then the perturbed system

$$(sI+s\Delta_0+(I+\Delta_n)L_n)\prod_{k=1}^{n-1}(sI+(I+\Delta_k)L_k)\,X = U_{\text{ref}},$$

where $U_{\text{ref}}, \Delta_k \in \mathcal{RH}_\infty$, achieves n^{th}-order consensus if

$$\|\Delta_k\|_{\mathcal{H}_\infty} < 1, \textit{ for all } k$$

and

$$\|\Delta_0\|_{\mathcal{H}_\infty} + \|\Delta_n\|_{\mathcal{H}_\infty} < 1.$$

Proof. First, note that we can construct $X(s) = \Xi_1(s)$, $s\Xi_k = -(I+\Delta_k)L_k\Xi_k + \Xi_{k+1}$ for $k = 1,\ldots,n-1$, and $s(I+\Delta_0)\Xi_n = -(I+\Delta_n)L_n\Xi_n + U_{\text{ref}}$. For Ξ_n we have exactly the first-order case of Theorem 2 and thus $\lim_{t\to\infty}\xi_n(t) = \alpha_n(t)\mathbf{1}_N$ if $\|\Delta_0\|_{\mathcal{H}_\infty} + \|\Delta_n\|_{\mathcal{H}_\infty} < 1$. Consider the following induction hypothesis: if $\Xi_{k+1}(s) = \mathbf{1}_N G_{k+1}(s) + H_{k+1}(s)$ where $H_{k+1}(s) \in \mathcal{RH}_\infty$, then $\Xi_k = \mathbf{1}_N G_k(s) + H_k(s)$ for some $H_k(s) \in \mathcal{RH}_\infty$. We then have

$$s\Xi_k = -(I+\Delta_k)L_k\Xi_k + \Xi_{k+1}$$

which can be represented by the block diagram Figure 3. Here, note that

$$L_k(sI+L_k)^{-1}\Xi_{k+1} = (sI+L_k)^{-1}L_k(H_{k+1}(s))$$

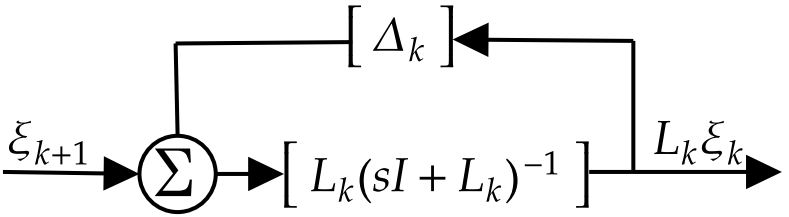

Figure 3. Block diagram illustrating the perturbation model of a general first-order consensus block which is used in the proof of Theorem 3.

and the potentially unstable term of Ξ_{k+1} can be ignored. Reusing a result from the previous proof we have $\|L_k(sI+L_k)^{-1}\|_{\mathcal{H}_\infty} = 1$ and therefore $L_k\Xi_k \in \mathcal{RH}_\infty$ if $\|\Delta_k\|_{\mathcal{H}_\infty} < 1$. Since the 0 eigenvalue of L_k is unique, it follows that $\Xi_k(s) = \mathbf{1}_N G_k(s) + H_k(s)$ with $H_k \in \mathcal{RH}_\infty$ which proves the induction hypothesis since we have already shown the base case $\Xi_n(s) = \mathbf{1}_N G_n(s) + H_n(s)$. It is left to prove that the system will reach n^{th}-order consensus. Note that $L_1X(s) = L_1\Xi_1(s)$ is stable and therefore we get through the final value theorem

$$\lim_{t\to\infty} L_1 x(t) = \lim_{s\to 0} sL_1\Xi_1(s) = 0.$$

Furthermore, we have for all k: $\lim_{s\to 0} sL_k\Xi_k(s) = 0$. This, combined with $s^2\Xi_k(s) = -(I+\Delta_k)sL_k\Xi_k(s) + s\Xi_{k+1}(s)$ shows that

$$\begin{aligned}\lim_{t\to\infty} L_{k+1}x^{(k)}(t) &= \lim_{s\to 0} s(s^k L_{k+1}X(s)) \\ &= \lim_{s\to 0} sL_{k+1}\Xi_{k+1}(s) = 0.\end{aligned}$$

Finally, since each L_k has a unique 0 eigenvalue with the corresponding eigenvector $\mathbf{1}_N$, we see that n^{th}-order consensus will be achieved □

This theorem shows that the n^{th}-order serial consensus is robust in its construction.

5. Examples

5.1 2^{nd}-Order Consensus on Circular Graph

Consider the directed cycle graph, which can be represented by the adjacency matrix

$$[W_{\text{cycle}}]_{i,j} = 1 \text{ iff } i - j = 1 \mod N.$$

The corresponding graph Laplacian L_{cycle} is a circulant matrix and therefore, the eigenvalues are known analytically. In particular, the eigenvalue with the second smallest real part is $\lambda_2(L_{\text{cycle}}) = 1 - \exp(2\pi\mathbf{i}/N) = 1 - \cos(2\pi/N) - \mathbf{i}\sin(2\pi/N)$.

For large N, this eigenvalue can be approximated with a first-order Taylor approximation, which yields $\lambda_2(L_c) \approx -\mathbf{i}2\pi/N$. This eigenvalue will cause problems when designing a controller using the conventional consensus. To see this, consider the closed loop dynamics

$$s^2 I + 2p_1 s L_{\text{cycle}} + p_0 L_{\text{cycle}} = U_{\text{ref}}.$$

The system can be diagonalized and, in particular, two of the poles are given by the equation $s^2 + 2p_1\lambda_2(L_{\text{cycle}}) + p_0\lambda_2(L_{\text{cycle}}) = 0$. In the case when p_0 and p_1 are designed independently of the network size N, then for sufficiently large N the roots can be approximated as

$$s_p = -p_1\lambda_2 \pm \sqrt{p_1^2\lambda_2^2 - p_0\lambda_2} \approx \pm(1+\mathbf{i})\sqrt{\frac{\pi p_0}{N}}.$$

Since one of these poles lies in the RHP, it follows that the closed loop system will become unstable when N is sufficiently large, for any (fixed) choice of p_0 and p_1.

For the serial consensus it is sufficient to check that all eigenvalues but the unique 0 eigenvalue of L_{cycle} lie in the RHP, or equivalently, if $0 < \text{Re}(\lambda_2(L_{\text{cycle}})) = 1 - \cos(2\pi/N)$ which is clearly true for any finite N. Alternatively, it is also sufficient to check that the underlying graph contains a connected spanning tree.

5.2 3rd-Order Consensus

It has been shown that the conventional consensus $x^{(n)} = -\sum_{k=0}^{n-1} L(\mathcal{G}_N)x^{(k)}$ cannot achieve scalable stability for any graph family $\{\mathcal{G}_N\}$ such that the corresponding graph Laplacian has an eigenvalue that decreases towards zero as the graph is growing, i.e. if $\lim_{N\to\infty} \text{Re}(\lambda_2(L(\mathcal{G}_N))) = 0$. At least, this is not possible with the conventional consensus control. However, for the serial consensus this is no longer a problem. The controller

$$U(s) = U_{\text{ref}} + \left(s^3 I - \prod_{k=1}^{3}(sI + L(\mathcal{G}_N))\right)X(s)$$

will achieve consensus as long as each of the underlying graphs $\{\mathcal{G}_N\}$ contains a connected spanning tree. To illustrate this, consider the graph defined by the adjacency matrix $W_{\text{path}} \in \mathbb{R}^{N\times N}$, defined as

$$[W_{\text{path}}]_{i,j} = \begin{cases} 1 & \text{if } |i-j| = 1 \text{ and } i \neq 1 \\ 0 & \text{else} \end{cases}.$$

This corresponds to a bidirectional path graph with a leader (Agent 1). Let L_{path} be the associated graph Laplacian. It is true that $\lim_{N\to\infty} \lambda_2(L_{\text{path}}) = 0$ and thus any conventional control design with L_{path} will eventually lead to an unstable closed loop. For this example, let the conventional control law be $u(t) = u_{\text{ref}}(t) - 6L_{\text{path}}\ddot{x} -$

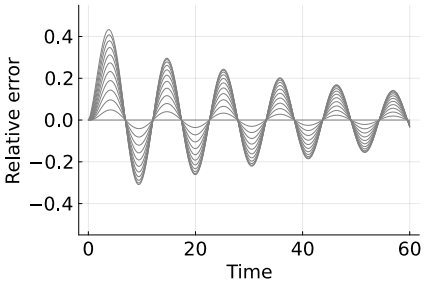

(a) Conventional consensus, $N = 12$.

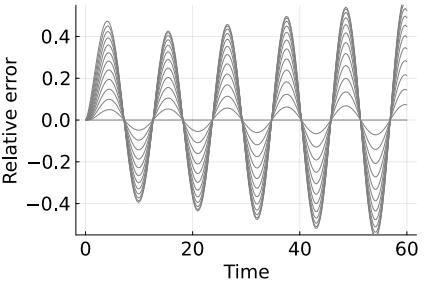

(b) Conventional consensus, $N = 13$.

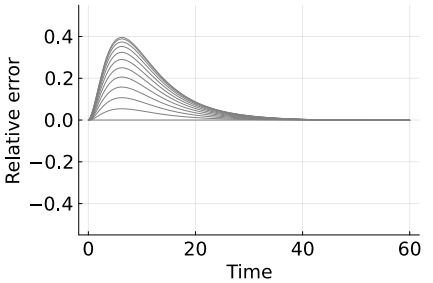

(c) Serial consensus, $N = 12$.

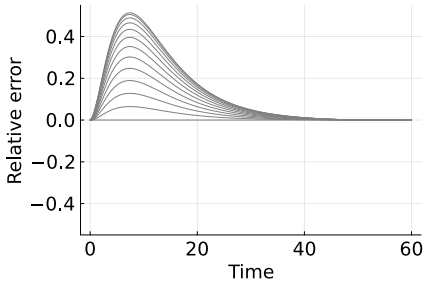

(d) Serial consensus, $N = 13$.

Figure 4. 3rd-order consensus in a chain of vehicles is considered. The plots show the inter-vehicle relative errors over time when the lead vehicle moves at constant acceleration. Panels (a) and (b) show that the addition of one agent destabilizes the closed loop for the conventional consensus. Panels (c) and (d) illustrate the fact that the serial consensus will remain stable under such agent additions.

$4L_{\text{path}}\dot{x} - 2L_{\text{path}}x$ and the serial consensus controller (5) be defined with the same graph Laplacians $L_k = 2kL_{\text{path}}$. The response to a constant acceleration of the leader is shown in Fig. 4. Here we see that the addition of a 13th agent to the system destabilizes the closed loop for the conventional consensus while the serial consensus only loses some performance.

5.3 Robustness of the 2nd-Order Serial Consensus.

Theorems 2 and 3 show that the serial consensus can be perturbed and still achieve n^{th}-order consensus. Now we want to illustrate what the block Δ_k can be. Consider the perturbed 2nd-order consensus system in Theorem 2. Writing out all terms we get

$$s^2(I+\Delta_2)X = U_{\text{ref}} - (s(2I+\Delta_1)LX + (I+\Delta_0)L^2X).$$

In this form, the Δ_2 block can represent potential model errors. While we might control a system modeled as N identical double-integrator systems, the reality may differ. This is obviously the case for vehicle platoons, which are often modeled as

chains of identical double integrators. Through our theorem we can for instance allow Δ_2 to be a diagonal transfer matrix with elements $[\Delta_2]_{i,i} = \frac{k_i}{T_i s+1}$ where $|k_i| < 1$ and $T_i > 0$ for all i. In this scenario, the closed loop system would remain stable despite the heterogeneous agents. The blocks Δ_1 and Δ_0 are also important. For instance, the signals $L^2 x(t)$ and $L\dot{x}(t)$ may not be directly measured, but instead estimated through linear filters. This could be thought of as unmodeled dynamics, which these blocks can capture.

If we focus on Theorem 3, then the perturbed model is

$$(s(\Delta_0 + I) + (\Delta_2 + I)L_2)\,(sI + (\Delta_1 + I)L_1)X = U_{\text{ref}}.$$

The theorem only asserts robustness for symmetrical graph Laplacians L_k. However, since each Δ_k can be a constant matrix, it is possible to construct new (asymmetric) graph Laplacians $L'_k = (I + \Delta_k)L_k$ by designing the Δ_k blocks.

6. Conclusions

This work has introduced the n^{th}-order serial consensus system, which serves as a natural generalization of the well-known consensus protocols. The stability of this system can be analyzed by considering n regular first-order consensus protocols. The controller proposed for achieving n^{th}-order serial consensus has been shown to be implementable using relative measurements confined to a local neighborhood of each agent and can therefore be considered a distributed control scheme. Robustness of the proposed system has also been analyzed. This has been addressed in terms of two different types of model perturbations. The analysis showed that the size, measured in the $\mathcal{H}_\infty$ norm, of the allowable uncertainties were independent of the number of agents.

Future and ongoing work will explore the performance of the serial consensus and its relation to string stability. It would also be interesting to consider an implementation where each agent employs an observer to compute their control action.

Appendix

Here we prove Lemmas 1, 2 which describe how the sparsity pattern of two matrices changes through addition and multiplication.

A Proof of Lemma 1

Proof. First, we have $\|A_1 + A_2\|_\infty \leq \|A_1\|_\infty + \|A_2\|_\infty \leq c_1 + c_2$ which follows from the triangle inequality.

For the second part we have $(A_1 + A_2)\mathbf{1}_N = 0 + 0 = 0$.

For the last part, WLOG, suppose that $q_1 \leq q_2 = \max(q_1, q_2)$. Since W is a positive matrix, we get

$$0 \leq \left(\sum_{k=0}^{q_1} W^k \right)_{i,j} \leq \left(\sum_{k=0}^{q_2} W^k \right)_{i,j}.$$

In particular, the following implication follows

$$\left(\sum_{k=0}^{q_2} W^k \right)_{i,j} = 0 \implies \left(\sum_{k=0}^{q_1} W^k \right)_{i,j} \implies [A_1 + A_2]_{i,j} = 0,$$

which concludes the proof. □

B Proof of Lemma 2

To prove the result on the product of two matrices, Lemma 2, we need the following three lemmas:

Lemma 3
Let $A, B \in \mathbb{C}^{N \times N}$ and define $\hat{A}_{i,j} = |A|_{i,j}$ and $\hat{B}_{i,j} = |B|_{i,j}$. If $(AB)_{i,j} \neq 0$ then $(\hat{A}\hat{B})_{i,j} \neq 0$.

Proof. Suppose the statement is false, i.e. $(\hat{A}\hat{B})_{i,j} = 0$ but $(AB)_{i,j} \neq 0$. Then we know that

$$(\hat{A}\hat{B})_{i,j} = \sum_{k=1}^{N} |A_{i,k}||B_{k,j}| = 0,$$

but this implies that at least one of $A_{i,k}$ and $B_{k,j}$ is equal to 0 for all k. But from this it follows that

$$(AB)_{i,j} = \sum_{k=1}^{N} A_{i,k} B_{k,j} = \sum_{k=1}^{N} 0 = 0.$$

This is a contradiction and concludes the proof. □

Lemma 4
Let $A, A_1, B, B_1 \in \mathbb{R}_+^{N \times N}$. If $(AB)_{i,j} \neq 0$ then $((A + A_1)(B + B_1))_{i,j} \neq 0$.

Proof. Expand the product to get

$$((A + A_1)(B + B_1))_{i,j} = (AB)_{i,j} + (AB_1)_{i,j} + (A_1 B)_{i,j} + (A_1 B_1)_{i,j} \geq (AB)_{i,j}$$

which followed from the fact that the product of 2 nonnegative matrices is also nonnegative. □

Lemma 5
Let $A, A_1, B, B_1 \in \mathbb{R}_+^{N\times N}$ be such that $A_{i,j} = 0$ if and only if $A_{1\,i,j} = 0$, and $B_{i,j} = 0$ if and only if $B_{1\,i,j} = 0$. Then, $(AB)_{i,j} = 0$ if and only if $(A_1B_1)_{i,j} = 0$.

Proof. The statement is clearly symmetrical and it is enough to prove sufficiency. Now, if $(AB)_{i,j} = 0$ then we know that

$$\sum_k A_{i,k}B_{k,j} = 0, \implies A_{i,k}B_{k,j} = 0, \ \forall k.$$

But this implies that either $A_{i,k} = 0$ or $B_{k,j} = 0$. In turn, this implies that either $(A_1)_{i,k} = 0$ or $(B_1)_{k,j} = 0$. And this leads to

$$(A_1B_1)_{i,j} = \sum_k (A_1)_{i,k}(B_1)_{k,j} = 0. \qquad \square$$

Now we can prove Lemma 2:

Proof. First, the gain can be bounded as $\|A_1A_2\|_\infty \leq \|A_1\|_\infty \|A_2\|_\infty \leq c_1c_2$ which followed from submultiplicity of the induced norm and from the definition of the sets.

For the second part we have $A_1A_2\mathbf{1}_N = A_1 0 = 0$.

For the last part we have to do slightly more. First replace each element in A_1 and A_2 with its absolute value and denote these B_1 and B_2. Now introduce two non-negative matrices C_1 and C_2 such that $B_1 + C_1 = 0 \iff \sum_{k=0}^{q_1} W^k$ and $B_2 + C_2 = 0 \iff \sum_{k=0}^{q_2} W^k$. Finally note that

$$(\sum_{k=0}^{q_1} W^k)(\sum_{j=0}^{q_2} W^j) = \sum_{k=0}^{q_1+q_2} w_k W^k$$

for some $w_k > 0$. By applying Lemma 5 two times we get that

$$\left[\sum_{k=0}^{q_1+q_2} W^k\right]_{i,j} = 0 \implies [(B_1 + C_1)(B_2 + C_2)]_{i,j} = 0.$$

Through Lemma 4 we get

$$[(B_1 + C_1)(B_2 + C_2)]_{i,j} = 0 \implies [B_1B_2]_{i,j} = 0.$$

And finally applying Lemma 3 results in

$$[B_1B_2]_{i,j} = 0 \implies [A_1A_2]_{i,j} = 0. \qquad \square$$

Acknowledgements

We want to thank Richard Pates for useful discussions regarding the robustness results.

The authors are with the Department of Automatic Control and the ELLIIT Strategic Research Area at Lund University, Lund, Sweden. This work was partially funded by Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation and the Swedish Research Council through Grant 2019-00691.

References

Bamieh, B., M. R. Jovanović, P. Mitra, and S. Patterson (2012). "Coherence in large-scale networks: Dimension-dependent limitations of local feedback". *IEEE Trans. Autom. Control* **57**:9, pp. 2235–2249.

Fax, J. A. and R. M. Murray (2004). "Information flow and cooperative control of vehicle formations". *IEEE Trans. Autom. Control* **49**:9, pp. 1465–1476.

Hansson, J. and E. Tegling (2023). "A closed-loop design for scalable high-order consensus". *Extended version*. arXiv:2304.12064.

Jadbabaie, A., J. Lin, and A. S. Morse (2003). "Coordination of groups of mobile autonomous agents using nearest neighbor rules". *IEEE Trans. Autom. Control* **48**:6, pp. 988–1001.

Jiang, F., L. Wang, and Y. Jia (2009). "Consensus in leaderless networks of high-order-integrator agents". *Am. Control Conf.*, pp. 4458–4463.

Olfati-Saber, R. and R. M. Murray (2004). "Consensus problems in networks of agents with switching topology and time-delays". *IEEE Trans. Autom. Control* **49**:9, pp. 1520–1533.

Pasqualetti, F., S. Zampieri, and F. Bullo (2014). "Controllability metrics, limitations and algorithms for complex networks". In: *Am. Control Conf.* Pp. 3287–3292.

Ren, W., K. L. Moore, and Y. Chen (2006). "High-order consensus algorithms in cooperative vehicle systems". In: *Proc. IEEE Int. Conf. Netw. Sens. Control*, pp. 457–462.

Ren, W., R. W. Beard, and E. M. Atkins (2007). "Information consensus in multivehicle cooperative control". *IEEE Control Syst. Mag.* **27**:2, pp. 71–82.

Seiler, P., A. Pant, and K. Hedrick (2004). "Disturbance propagation in vehicle strings". *IEEE Trans. Autom. Control* **49**:10, pp. 1835–1842.

Siami, M. and N. Motee (2016). "Fundamental limits and tradeoffs on disturbance propagation in large-scale dynamical networks". *IEEE Trans. Autom. Control* **61**:12, pp. 4055–4062.

Stüdli, S., M. M. Seron, and R. H. Middleton (2017). "Vehicular platoons in cyclic interconnections with constant inter-vehicle spacing". *IFAC-PapersOnLine* **50**:1, pp. 2511–2516.

Swaroop, D. and J. K. Hedrick (1996). "String stability of interconnected systems". *IEEE Trans. Autom. Control* **41**:3, pp. 349–357.

Tegling, E., B. Bamieh, and H. Sandberg (2023). "Scale fragilities in localized consensus dynamics". *Automatica* **153**, p. 111046.

Zhou, K. and J. C. Doyle (1998). *Essentials of robust control*. Vol. 104. Prentice hall Upper Saddle River, NJ.

Paper II

Closed-Loop Design for Scalable Performance of Vehicular Formations

Jonas Hansson **Emma Tegling**

Abstract

This paper presents a novel control design for vehicular formations as an alternative to alignment through conventional consensus protocols for second-order systems. The design is motivated by the closed-loop system, which we construct as first-order systems connected in series, and is therefore called *serial consensus*. The serial consensus design will guarantee the stability of the closed-loop system under the minimum requirement of the underlying communication graph containing a directed spanning tree—which is not generally true for conventional consensus. As our main result, we show that the serial consensus design gives bounds on the worst-case transient behavior of the formation, which is independent of the number of vehicles and the underlying graph structure. In particular, this shows that the serial consensus design guarantees string stability of the formation and is, therefore, suitable for directed formations and communication topologies. We show that serial consensus can be implemented through message passing or measurements to neighbors at most two hops away. We illustrate our results through numerical examples.

1. Introduction

Network systems emerge in a wide range of applications, and engineered networks are, in many cases, becoming increasingly large-scale and complex. Examples include smart power grids, sensor networks, traffic, and multi-robot networks, where coordinating a multitude of interconnected subsystems or agents is a crucial control problem. The prototypical coordination problem that leads to distributed consensus dynamics was studied early on by [Fax and Murray, 2004; Olfati-Saber and Murray, 2004; Jadbabaie et al., 2003], and the dynamic behaviors of this and related problems have since been the subject of much research. The literature has made clear that on large scales, consensus-type networks often exhibit poor dynamic behaviors, for example in terms of controllability [Pasqualetti et al., 2014], performance and coherence [Bamieh et al., 2012; Siami and Motee, 2016; Tegling et al., 2019], disturbance propagation [Swaroop and Hedrick, 1996; Seiler et al., 2004; Besselink and Knorn, 2018] and even instability [Tegling et al., 2023]. Motivated by these issues, our work proposes an alternative consensus control design with fundamentally improved scalability properties.

We consider a simple yet classical vehicular formation control problem in which each vehicle of the formation is modeled as a double integrator whose controller relies on relative state measurements between neighboring vehicles. In the one-dimensional case (longitudinal control), this approach can be compactly written on the conventional second-order consensus form

$$\ddot{x}(t) = u(x,t) = -L_{\text{vel}}\dot{x}(t) - L_{\text{pos}}x(t) + u_{\text{ref}}(t). \tag{1}$$

Here, x is a vector that represents the position of each vehicle, u is a vector of the control inputs, L_{vel} and L_{pos} are graph Laplacians respectively representing the velocity and positional feedback, and u_{ref} is a reference control signal. This work will focus on the essentials of vehicular formation control and the connection between transient performance and communication structure, for which this simple model is suitable. However, the literature contains several more advanced and accurate formation models, like the third-order model in [Wijnbergen et al., 2021] and the fourth-order one in [Dolk et al., 2017].

Over the years, various assumptions on the feedback structure, here captured by the two graph Laplacians in (1), have been considered. They were assumed to be proportional to each other in the early work [Ren and Atkins, 2007], and in more recent analyses [Patterson and Bamieh, 2014; Tegling et al., 2023]. In this case, when the Laplacians capture relative and localized feedback, there are at least three problems with the design (1). First, stability is not guaranteed for all graph Laplacians. For example, the system may be unstable if the Laplacian corresponds to a directed cycle graph, see e.g. [Stüdli et al., 2017b; Cantos et al., 2016]. Second, in directed vehicle strings, small errors may amplify throughout the formation and lead to so-called string instability—a topic thoroughly surveyed in [Stüdli et al., 2017a; Feng et al., 2019]. Third, in the case of undirected vehicle

formations, the convergence rate of the formation may scale poorly (as $O(1/N^2)$) [Barooah et al., 2009]. This paper's primary focus is error amplification avoidance, under the restriction of only allowing the vehicles knowledge of their neighborhood and none of the formation size. The bounded transient, regardless of formation size, together with the locality restrictions, is what we refer to as having *scalable performance*. Although our work is mainly theoretical, the problem of achieving scalable performance is also a practical one. In [Gunter et al., 2021], they noted that modern, commercial, adaptive cruise control systems for cars lack scalable performance and thus fail to maintain a formation.

The question of suitable communication structures for the feedback law (1) is not new. Indeed, a systematic study [Herman et al., 2017] noted that using different Laplacians for the position and velocity dynamics in (1) may drastically improve the performance of the formation. Specifically, symmetric position feedback and asymmetric velocity feedback were proposed, i.e., $L_{\text{pos}}^T = L_{\text{pos}}$, which may be an idealization. However, the stability proof relies on L_{pos} being symmetric. In general, these systems are not straightforward to analyze in terms of the underlying topological properties, especially when considering scalability, that is, the growth of the network. Several analytic results can, however, be derived under assumptions of spatial invariance, that is, identical agents using the same control and interaction laws [Bamieh et al., 2008; Bamieh et al., 2012].

In this paper, we propose a new controller for the vehicle formation: $u(x,t) = -(L_1+L_2)\dot{x}(t) - L_2L_1x(t) + u_{\text{ref}}(t)$, where L_1 and L_2 are graph Laplacians (their product, however, is in general no longer a Laplacian). The controller is designed to give a particular closed-loop system that we term the second-order *serial consensus system*. The name, as well as the reason for choosing this particular control structure, is easiest seen by considering the closed-loop system in the Laplace domain:

$$(sI+L_2)(sI+L_1)X(s) = U_{\text{ref}}(s).$$

The closed-loop system has the same dynamics as two conventional first-order consensus systems put in a series, but unlike first-order consensus, we now control the second-order dynamics of the vehicles. As for the classical first-order consensus protocol, the consensus equilibrium will be stable if the graphs underlying L_1 and L_2 each contain a directed spanning tree. Since this is usually assumed true, this directly addresses the earlier mentioned instability problem in conventional second-order consensus. However, this paper's main result concerns the serial consensus's performance.

It turns out that the serial consensus controller can guarantee a strong notion of (generalized) string stability of the formation. Specifically, given any combination of relative errors $e_{\text{p}}(t) = Lx(t)$, that is defined through a graph Laplacian L, and velocity deviations $e_{\text{v}}(t) = \dot{x}(t) - \mathbf{1}v_{\text{ref}}$, we can give bounds on the following form:

$$\sup_{t\geq 0} \left\| \begin{bmatrix} e_{\text{p}}(t) \\ e_{\text{v}}(t) \end{bmatrix} \right\|_\infty \leq \alpha \left\| \begin{bmatrix} e_{\text{p}}(0) \\ e_{\text{v}}(0) \end{bmatrix} \right\|_\infty .$$

Here, α is a constant independent of the number of vehicles and the underlying graph structure (that is, it need not be a string, though it holds for the directed string in particular). Our result is significant because it is stated in terms of the supremum of the ℓ^∞-norm, and thus directly bounds the time evolution of the transient. Results of this type have been suggested to be more suitable for large vehicle formations [Feintuch and Francis, 2012; Stüdli et al., 2017a], mainly to ensure scalability [Besselink and Knorn, 2018], although typically hard to derive and therefore analyzed through proxies. The result implies that our design addresses the earlier mentioned problems of the general conventional consensus protocol; under mild conditions, the closed loop will be stable, and it can be designed to achieve string stability. Since directed graphs are allowed, it is also possible to avoid the poor scaling of the algebraic connectivity and thus achieve a faster convergence rate [Barooah et al., 2009; Hao and Barooah, 2012].

The cost for this advantage is sometimes a requirement of an additional communication step, either through physical measurement or signaling. Such additional signaling has been proposed in a vehicular platooning context, e.g., [Darbha et al., 2019]. We argue, however, that our structure gives more significant benefits with a smaller communications overhead.

The remainder of the paper is organized as follows. Section 2 introduces the problem setup and the notation used throughout the paper. Here, the serial consensus system is also defined together with some fundamental properties. In Section 3, we present our performance results as a theorem. The results are illustrated in Section 5 through numerical examples. Finally, we conclude the paper in Section 6.

2. Problem Setup

2.1 Definitions and network model

Let $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ denote a graph of size $N = |\mathcal{V}|$ with the edge set $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$. The graph can be equivalently represented by the weighted adjacency matrix $W \in \mathbb{R}^{N \times N}$ where $w_{i,j} > 0 \iff (j,i) \in \mathcal{E}$. The graph is called *undirected* if $W = W^T$. The graph contains a *directed spanning tree* if, for some $i \in \mathcal{V}$, there is a path from i to any other vertex $j \in \mathcal{V}$.

The weighted graph Laplacian L associated to the graph is defined as

$$[L]_{i,j} = \begin{cases} -w_{i,j}, & \text{if } i \neq j \\ \sum_{k \neq i} w_{i,k}, & \text{if } i = j \end{cases}. \tag{2}$$

Under the condition that that the graph generating the graph Laplacian contains a directed spanning tree, L will have a simple and unique eigenvalue at 0 and the remaining eigenvalues will lie strictly in the right half plane (RHP). We will refer to any $N \times N$ matrix that satisfies (2) for some set of non-negative weights $w_{i,j}$ as a graph Laplacian.

In this work, we also consider networks with a growing number of nodes. With $\mathcal{G}_N = (\mathcal{V}_N, \mathcal{E}_N)$ we denote a graph in a family $\{\mathcal{G}_N\}$, where N is the size of the growing network.

We will denote the space of all proper, real rational, and stable transfer matrices $\mathcal{RH}_\infty$ and denote the $\mathcal{H}_\infty$-norm as $\|\cdot\|_{\mathcal{H}_\infty}$ following the notation in [Zhou and Doyle, 1998]. By $\|\cdot\|_\infty$, we denote the standard vector norm and its corresponding induced matrix norm i.e. $\|z\|_\infty = \max_k |z_k|$, where $z \in \mathbb{C}^N$ and with $\|M\|_\infty = \sup_{\|x\|_\infty = 1} \|Mx\|_\infty$, for $M \in \mathbb{C}^{N\times N}$.

2.2 Vehicle formation model

Consider a simple vehicle formation which consists of N identical double-integrator systems, i.e.

$$\frac{\mathrm{d}^2 x_i(t)}{\mathrm{d}t^2} = u_i(x,t), \quad i = 1,\ldots,N, \tag{3}$$

where $x_i(t) \in \mathbb{R}$. The aim is to coordinate the vehicles to keep a fixed spacing and common velocity. This goal is related to the problem of achieving *second-order consensus* as defined below.

Definition 1—Second-order consensus
The vehicle formation (3) is said to achieve second-order consensus if

$$\lim_{t\to\infty} \left| \frac{\mathrm{d}x_i(t)}{\mathrm{d}t} - \frac{\mathrm{d}x_j(t)}{\mathrm{d}t} \right| = 0 \text{ and } \lim_{t\to\infty} |x_i(t) - x_j(t)| = 0$$

for all $i, j \in \mathcal{V}$.

With our control structure, the desired, fixed intervehicle distances can be set to zero for analysis purposes without loss of generality. This will be clarified in Remark 2.

Remark 1
In this work, we only consider a scalar state, that is, longitudinal control. Our approach can be extended to higher spatial dimensions (see e.g. [Oh et al., 2015] for a survey of approaches), but we omit it here to keep the notation simple.

2.3 Control structure

In this work, we consider linear state feedback controllers of the system (3). Such controllers can be written as

$$u(t) = u_{\mathrm{ref}}(t) - A_{\mathrm{vel}}\dot{x}(t) - A_{\mathrm{pos}}x(t), \tag{4}$$

where $x = (x_1, x_2, \ldots, x_N)^T \in \mathbb{R}^N$, $u_{\mathrm{ref}}(t) \in \mathbb{R}^N$ is a feedforward term, and $A_{\mathrm{pos}}, A_{\mathrm{vel}} \in \mathbb{R}^{N\times N}$ are constant feedback matrices for the position and velocity respectively. In the distributed coordination problem, the controller is further restricted to

i) only use *relative* feedback;

ii) have a bounded gain;

iii) only depend on the local neighborhood of each agent.

These restrictions are captured by considering controllers in the following class.

Definition 2—q-step implementable relative feedback
The relative state feedback $u = Ax$ is q-step implementable with respect to the adjacency matrix W and gain $c > 0$ if $A \in \mathcal{A}^q(W,c)$, where

$$\mathcal{A}^q(W,c) = \left\{ A \;\middle|\; \begin{array}{c} \left[\sum_{k=0}^{q} W^k\right]_{i,j} = 0 \implies A_{i,j} = 0, \\ A\mathbf{1} = 0,\ \|A\|_\infty \le c \end{array} \right\}.$$

Clearly, the sum of two q-step implementable controllers is also q-step implementable, so if both $A_{\text{pos}}, A_{\text{vel}} \in \mathcal{A}^q(W,c)$ then the combined controller in (4) will also be q-step implementable. To clarify the concept of q-step implementability, consider the following example.

Example 1
Consider a vehicle string where each agent can measure the distance to its two neighboring vehicles. This structure can be represented by the adjacency matrix W such that $[W]_{i,j} = 1 \iff |i-j| = 1$ and $W_{i,j} = 0$ otherwise. Then the sparsity constraint of $A \in \mathcal{A}^q(W,c)$ corresponds to the requirement that

$$|i-j| > q \implies [A]_{i,j} = 0,$$

i.e., only relative measurements up to the q nearest neighbors are used. One choice of a 1-step implementable matrix ($A \in \mathcal{A}^1(W,c)$) is the graph Laplacian for an undirected path graph $A = L_{\text{undir-path}}$ while an example of a 2-step implementable matrix is $A = L^2_{\text{undir-path}}$.

In general, if the adjacency matrix W captures a physical network, then a controller $u = Ax$ with $A \in A^q(W,c)$ means u_i only requires signals from Agent i's q-hop neighborhood. This is readily proven; we refer the reader to [Hansson and Tegling, 2023].

A controller widely applied for vehicle formations in the literature is what we will call the *conventional consensus* controller. In this case, both $A_{\text{pos}} = L_{\text{pos}}$ and $A_{\text{vel}} = L_{\text{vel}}$ are chosen to be graph Laplacians and this results in the controller being a 1-step implementable relative-feedback controller. The closed-loop dynamics with this controller are

$$\ddot{x} = -L_{\text{vel}}\dot{x} - L_{\text{pos}}x + u_{\text{ref}}. \tag{5}$$

But, this is not the only way to implement a controller satisfying the desired structure. We next propose our alternative approach.

Remark 2
The analysis of a formation with position offsets can be made on the translated states $\tilde{x} = x - d - tv_{\text{ref}}\mathbf{1}$ *where* $d \in \mathbb{R}^N$ *is a vector of desired offsets and* v_{ref} *a desired velocity. If the reference control signal is chosen to be* $u_{\text{ref}} = \tilde{u}_{\text{ref}} + A_{\text{pos}}d$ *then the closed-loop dynamics in the new states becomes*

$$\ddot{\tilde{x}} = \tilde{u}_{\text{ref}} - A_{\text{vel}}\dot{\tilde{x}} - A_{\text{pos}}\tilde{x}$$

where the property $A_{\text{vel}}\mathbf{1} = A_{\text{pos}}\mathbf{1} = 0$ *was used. Thus, the dynamics around any possible offset will be equivalent to the dynamics of* x *when* $d = 0$ *and* $v_{\text{ref}} = 0$.

2.4 A Novel Design: Serial Consensus

We propose a control design that achieves a desired closed loop to address the stability and performance of interconnected double-integrator systems. Due to its structure, we call it the second-order serial consensus system.

Definition 3—Second-order serial consensus system
Let L_1 *and* L_2 *be weighted and directed graph Laplacians. The second-order serial consensus system is then*

$$(sI + L_2)(sI + L_1)X(s) = U_{\text{ref}}(s). \tag{6}$$

The system (3) achieves the serial consensus system through the control design

$$u(x,t) = u_{\text{ref}}(t) - (L_2 + L_1)\dot{x}(t) - L_2L_1x(t). \tag{7}$$

Remark 3
The new controller (7) is similar to the corresponding controller used in (5). This can be seen through the matching L_{vel} *to* $L_2 + L_1$ *and* L_{pos} *to* L_2L_1. *However, the product* L_2L_1 *will only be a graph Laplacian in special cases. Thus, the serial consensus controller adds a new set of controllers not previously considered in the literature.*

When analyzing the serial consensus system (6), we will use the following assumption on the graph structure.

Assumption 1—Directed spanning tree
The graphs underlying the graph Laplacians L_1 *and* L_2 *each contain a directed spanning tree.*

A convenient state-space representation of the serial consensus system is

$$\begin{bmatrix}\dot{\xi}_1\\ \dot{\xi}_2\end{bmatrix} = \underbrace{\begin{bmatrix}-L_1 & I\\ 0 & -L_2\end{bmatrix}}_{A}\begin{bmatrix}\xi_1\\ \xi_2\end{bmatrix} + \begin{bmatrix}0\\ u_{\text{ref}}\end{bmatrix}. \tag{8}$$

This can be transformed back to x and $\dot{x}$ through the linear transformations

$$\begin{bmatrix} x \\ \dot{x} \end{bmatrix} = \begin{bmatrix} I & 0 \\ -L_1 & I \end{bmatrix} \begin{bmatrix} \xi_1 \\ \xi_2 \end{bmatrix} \text{ and } \begin{bmatrix} \xi_1 \\ \xi_2 \end{bmatrix} = \begin{bmatrix} I & 0 \\ L_1 & I \end{bmatrix} \begin{bmatrix} x \\ \dot{x} \end{bmatrix}.$$

The following theorem explains some benefits of considering the serial consensus.

Theorem 1
The second-order serial consensus system (6) under Assumption 1 and with $U_{\text{ref}} \in \mathcal{RH}_\infty$ has the properties:

i) its poles are given by the union of the eigenvalues of $-L_1$ and $-L_2$;

ii) its solution achieves second-order consensus.

The proof, a version of which appeared in [Hansson and Tegling, 2023], is presented in Appendix B. Per Theorem 1, the serial consensus system has a stable consensus equilibrium, and since this holds independently of the number of agents, it is also scalably stable. For the conventional consensus, a theorem like this does not exist since using, e.g., a Laplacian corresponding to a directed cycle can result in an unstable closed loop as noted in [Stüdli et al., 2017b; Tegling et al., 2023]. Serial consensus, on the other hand, can also be robustly stable; see [Hansson and Tegling, 2023] for the robustness criteria.

At worst, the serial consensus controller is 2-step implementable as per the following result.

Proposition 1
Consider the second-order serial consensus controller (7). If $L_1, L_2 \in \mathcal{A}^1(W,c)$ for an adjacency matrix W and a constant c, then the controller is a 2-step implementable relative feedback controller with respect to W and gain $c' = \max\{2c, c^2\}$.

The proof is in Appendix C, a version of what appeared in [Hansson and Tegling, 2023]. The implementation of the serial consensus will be further discussed after our main results.

2.5 Performance criterion

Motivated by the vehicle formation setting, we introduce the following two errors. First, the relative position error is defined by

$$e_{\text{p}}(t) := d + Lx(t),$$

where the measurement graph $\mathcal{G}_N$ underlying the graph Laplacian L has size N and $d \in \mathbb{R}^N$ is a vector of desired offsets. Second, denote the velocity deviation as

$$e_{\text{v}}(t) := \dot{x}(t) - v_{\text{ref}}\mathbf{1},$$

with $v_{\text{ref}} \in \mathbb{R}$ being the desired vehicle velocity.

The error e_{p} represents the local relative position errors. This error needs to remain small to prevent vehicle collisions. Meanwhile, e_{v} represents the deviation from the desired velocity. This error must be small to ensure that speed limits are respected. This should also remain true as the network grows; the errors should be independent of the number of agents N.

Definition 4—Scalable performance
A formation controller defined over a growing family of graphs $\{\mathcal{G}_N\}$ that ensures

$$\sup_{t\geq 0} \left\| \begin{bmatrix} e_{\text{p}}(t) \\ e_{\text{v}}(t) \end{bmatrix} \right\|_\infty \leq \alpha \left\| \begin{bmatrix} e_{\text{p}}(0) \\ e_{\text{v}}(0) \end{bmatrix} \right\|_\infty, \tag{9}$$

where α is fixed and independent of the number of agents N, is said to achieve scalable performance.

Here, the choice of norm is essential since the worst-case behavior gets bounded by the initial maximum deviation. We remark that we only consider the initial value response. While the disturbance amplification scenario also requires careful analysis, as discussed in [Besselink and Knorn, 2018], we leave it outside the scope of the present study.

3. Main Results

Here, we show that position and velocity errors can be kept small throughout the transient phase, regardless of network size. This is achieved using a serial consensus controller that uses measurements based on the underlying network graph. The result is summarized in the following theorem.

Theorem 2
Let $e_{\text{p}} = Lx$ where L is a graph Laplacian. If the system $\ddot{x}(t) = u(t)$ is controlled with

$$u(t) = -(p_1 + p_2)L\dot{x}(t) - p_1 p_2 L^2 x(t), \tag{10}$$

where $p_1, p_2 > 0$ and $p_1 \neq p_2$, then the resulting serial consensus system achieves scalable performance with

$$\alpha = \frac{1}{|p_1 - p_2|}\left(p_1 + p_2 + \max\{2, 2p_1 p_2\}\right).$$

Proof. The serial consensus system can be rewritten as

$$\begin{bmatrix} \dot{\xi}_1 \\ \dot{\xi}_2 \end{bmatrix} = \begin{bmatrix} -p_1 L & I \\ 0 & -p_2 L \end{bmatrix} \begin{bmatrix} \xi_1 \\ \xi_2 \end{bmatrix} + \begin{bmatrix} 0 \\ u_{\text{ref}} \end{bmatrix}.$$

Here $\xi_1 = x$ and $\xi_2 = \dot{x} + p_1 L x$. Since $u_{\text{ref}} = 0$, the initial value problem can be solved directly. This evaluates to

$$\xi_1(t) = e^{-p_1 L t}\xi_1(0) + e^{-p_1 L t}\int_0^t e^{p_1 L\tau} e^{-p_2 L\tau}\mathrm{d}\tau\xi_2(0)$$
$$\xi_2(t) = e^{-p_2 L t}\xi_2(0).$$

Since $p_1 L$ and $p_2 L$ obviously commute, it follows that $e^{p_1 L\tau}e^{-p_2 L\tau} = e^{(p_1-p_2)L\tau}$. By pre-multiplying the first equation with $(p_1 - p_2)L$ and using the property that L commutes with $e^{-p_1 L t}$ we get the integrand $(p_1 - p_2)Le^{(p_1-p_2)L\tau} = \frac{\mathrm{d}}{\mathrm{d}t}(e^{(p_1-p_2)L\tau})$. Finally, by applying the Fundamental Theorem of Calculus, the equation can be simplified to

$$L\xi_1(t) = e^{-p_1 L t}L\xi_1(0) + \frac{(e^{-p_2 L t} - e^{-p_1 L t})}{p_1 - p_2}\xi_2(0).$$

Now, inserting that $e_{\text{p}}(t) = L\xi_1(t), \dot{x}(t) = \xi_2(t) - p_1 e_{\text{p}}(t)$, and $e_{\text{v}}(t) = \dot{x} - v_{\text{ref}}\mathbf{1}$ yields

$$e_{\text{p}}(t) = e^{-p_1 L t}e_{\text{p}}(0) + \frac{e^{-p_2 L t} - e^{-p_1 L t}}{p_1 - p_2}(e_{\text{v}}(t) + v_{\text{ref}}\mathbf{1} + p_1 e_{\text{p}}(0)).$$

Next we note that the relation $e^{-Lt}\mathbf{1} = \mathbf{1}$ holds for any graph Laplacian L. In particular, this implies that $(e^{-p_2 L t} - e^{-p_1 L t})\mathbf{1} = 0$. After some simplifications, this leads to

$$e_{\text{p}}(t) = \frac{p_1 e^{-p_2 L t} - p_2 e^{-p_1 L t}}{p_1 - p_2}e_{\text{p}}(0) + \frac{e^{-p_2 L t} - e^{-p_1 L t}}{p_1 - p_2}e_{\text{v}}(0)$$

$$e_{\text{v}}(t) = -p_1 e_{\text{p}}(t) + e^{-L_2 t}(e_{\text{v}}(0) + p_1 e_{\text{p}}(0))$$
$$= \frac{p_1 p_2\left(e^{-p_1 L t} - e^{-p_2 L t}\right)}{p_1 - p_2}e_{\text{p}}(0) + \frac{p_1 e^{-p_1 L t} - p_2 e^{-p_2 L t}}{p_1 - p_2}e_{\text{v}}(0).$$

Taking the induced norm $\|\cdot\|_\infty$ and then applying the triangle inequality yields the following two bounds:

$$\|e_{\text{p}}(t)\|_\infty \leq \frac{p_1 + p_2}{|p_1 - p_2|}\|e_{\text{p}}(0)\|_\infty + \frac{2}{|p_1 - p_2|}\|e_{\text{v}}(0)\|_\infty$$

and

$$\|e_{\text{v}}(t)\|_\infty \leq \frac{2p_1 p_2}{|p_1 - p_2|}\|e_{\text{p}}(0)\|_\infty + \frac{p_1 + p_2}{|p_1 - p_2|}\|e_{\text{v}}(0)\|_\infty.$$

Finally, we have that

$$\|e_{\text{p}}\|_\infty, \|e_{\text{v}}\|_\infty \leq \left\|\begin{bmatrix} e_{\text{p}} \\ e_{\text{v}} \end{bmatrix}\right\|_\infty = \max\{\|e_{\text{p}}\|_\infty, \|e_{\text{v}}\|_\infty\}$$

Combining these facts yields

$$\left\| \begin{bmatrix} e_{\mathrm{p}}(t) \\ e_{\mathrm{v}}(t) \end{bmatrix} \right\|_{\infty} \leq \frac{p_1 + p_2 + 2\max\{1, p_1 p_2\}}{|p_1 - p_2|} \left\| \begin{bmatrix} e_{\mathrm{p}}(0) \\ e_{\mathrm{v}}(0) \end{bmatrix} \right\|_{\infty},$$

which is the definition of scalable performance with α as in the theorem statement.□

Remark 4
Unlike most consensus results, Theorem 2 does not need Assumption 1, requiring an underlying spanning tree. This is possible since the relative errors between disconnected components would be unobservable because the same graph Laplacian is used for the control and performance metric.

Remark 5
Theorem 2 excludes the case $p_1 = p_2$. It turns out that, in this case, there is no finite upper bound on the transient. For instance, when considering the directed path graph, the transient scales at least proportionally to $\sqrt{N}$; we show this in Appendix A.

In the limit when p_1 approaches infinity and p_2 approaches 0 (or vice versa), then the theoretically optimal bound of $\alpha = 1$ is retrieved. For instance, let $p_1 = \gamma$ and $p_2 = 1/\gamma$. Then, for $\gamma > 1$ we get

$$\lim_{\gamma \to \infty} \alpha(\gamma) = \lim_{\gamma \to \infty} \frac{\gamma + 1/\gamma + 2}{\gamma - 1/\gamma} = 1.$$

In this case, the dynamics are essentially reduced to a first-order consensus system, which may be desirable. This would, however, require an unbounded gain. We also note that the assumption in Theorem 2 can be relaxed: instead of requiring $L_1 = p_1 L$ and $L_2 = p_2 L$ to be proportional, it is sufficient for L_1 and L_2 to commute. Under this weaker condition, a similar performance result can be established for the positional errors defined as $e_p = (L_1 - L_2)x$, rather than the current $e_p = Lx$.

4. Implementation

The serial consensus protocol has now been shown to be scalable in both stability and performance. This also holds for the implementation. In the vehicle formation setting, we desire a controller with a finite gain and decentralized implementation that uses few measurements.

As seen from (7), the controller must implement the graphs underlying $L_1 + L_2$ and $L_2 L_1$. While this requires carefully designed signaling, it is still easy to implement as a controller with finite gain and localized measurements.

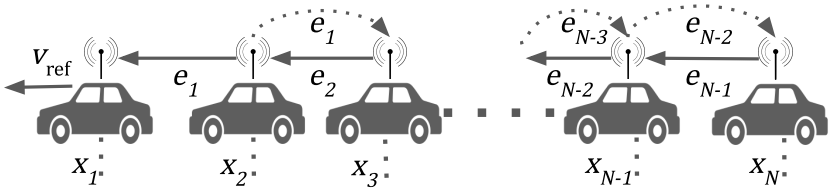

Figure 1. Vehicle platoon with directed measurements and message passing to implement serial consensus.

4.1 Message passing

One way to implement serial consensus is through message passing. The reason can be explained through the control law, which, as we recall, is

$$u = -(L_2 + L_1)\dot{x} - L_2 L_1 x + u_{\text{ref}}.$$

The velocity feedback will be implementable using only relative measurements from each agent, provided that both L_1 and L_2 are localized graph Laplacians. Message passing is needed for positional feedback. This can be implemented if each agent registers their relative error $e_i = [L_1 x]_i$. Then, if each agent can access their out-neighbor's error, the control signal can be calculated using $L_2 L_1 x = L_2 e$. In the case of a vehicle platoon, the relative distance to the first neighbor can be measured using radar. However, the relative distance to the second neighbor requires an additional signaling layer to implement message passing. If such signaling can be had with the nearest neighbor, it is implementable in a platoon. This idea of signaling is illustrated in Figure 1.

4.2 Extended measurements

Using an additional communication step is not the only way to implement the serial consensus. Instead, $L_2 L_1$ can be implemented through direct measurements. Indeed, with careful design, it is possible to choose L_2, L_1 so that their product $(L_2 L_1) \in \mathcal{A}^1(W, c)$ and is thus implementable using only relative measurements with immediate neighbors. The following example illustrates this case.

Example 2
Let $W_{\text{undir}-\text{path}}$ *correspond to the undirected path graph, i.e.*

$$(W_{\text{undir}-\text{path}})_{i,j} = 1 \iff |i-j| = 1.$$

Furthermore, let $L_{\text{ahead}-\text{path}}$ *and* $L_{\text{behind}-\text{path}}$ *correspond to the look-ahead and look-behind path graphs, respectively:*

$$L_{\text{ahead}-\text{path}} = \begin{bmatrix} 0 & 0 & & \\ -1 & 1 & & \\ & \ddots & \ddots & \\ & & -1 & 1 \end{bmatrix} \in \mathcal{A}^1(W_{\text{undir}-\text{path}}, 2) \tag{11}$$

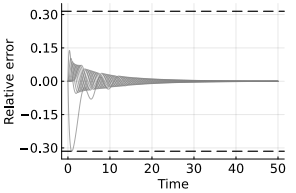

(a) Directed cycle, conventional consensus with $N = 10$ and transient $\sup_t \|e_\mathrm{p}\|_\infty = 0.137$.

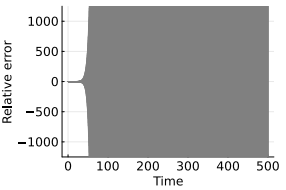

(b) Directed cycle, conventional consensus with $N = 100$ and transient $\sup_t \|e_\mathrm{p}\|_\infty = \infty$.

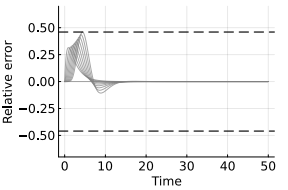

(c) Directed path, conventional consensus with $N = 10$ and transient $\sup_t \|e_\mathrm{p}\|_\infty = 0.460$.

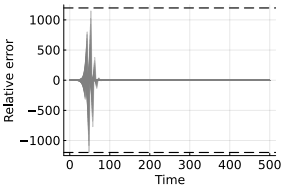

(d) Directed path, conventional consensus with $N = 100$ and transient $\sup_t \|e_\mathrm{p}\|_\infty = 1200$.

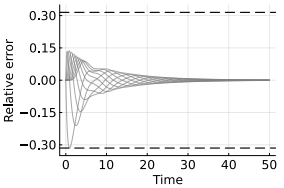

(e) Directed cycle, serial consensus with $N = 10$ and transient $\sup_t \|e_\mathrm{p}\|_\infty = 0.211$.

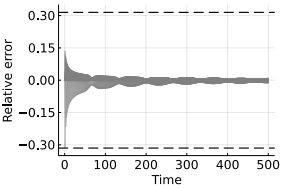

(f) Directed cycle, serial consensus with $N = 100$ and transient $\sup_t \|e_\mathrm{p}\|_\infty = 0.217$.

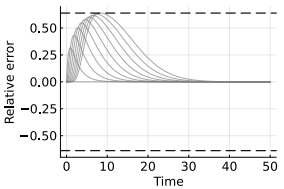

(g) Directed path, serial consensus with $N = 10$ and transient $\sup_t \|e_\mathrm{p}\|_\infty = 0.639$.

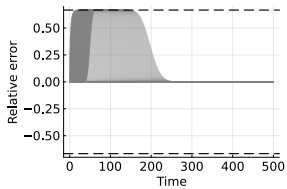

(h) Directed path, serial consensus with $N = 100$ and transient $\sup_t \|e_\mathrm{p}\|_\infty = 0.667$.

Figure 2. Simulation of the initial value response to $x = 0$, $\dot{x}_{i\neq 1}(0) = 0$, and $\dot{x}_1(0) = 1$. For the serial consensus, $p_1 = 2$ and $p_2 = 0.5$ were used, and for the conventional, $r_\mathrm{v} = 2.5$ and $r_\mathrm{p} = 1$ were used. Both different graph structures and the number of vehicles N were tested. For each plot, the inter-vehicle distances are shown. The conventional consensus system can degrade with increasing vehicles N, while the serial consensus displays scalable stability and performance.

and

$$L_{\text{behind}-\text{path}} = \begin{bmatrix} 1 & -1 & & \\ & \ddots & \ddots & \\ & & 1 & -1 \\ & & 0 & 0 \end{bmatrix} \in \mathcal{A}^1(W_{\text{undir}-\text{path}}, 2). \tag{12}$$

Then, the product of these two matrices will be

$$L_{\text{behind}-\text{path}}L_{\text{ahead}-\text{path}} = \begin{bmatrix} 1 & -1 & & & \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ & & & 0 & 0 \end{bmatrix}.$$

Since the product only requires information from the neighboring states, it holds that

$$L_{\text{behind}-\text{path}}L_{\text{ahead}-\text{path}} \in \mathcal{A}^1(W_{\text{undir}-\text{path}}, 4),$$

and the same holds true for $(L_{\text{behind}-\text{path}} + L_{\text{ahead}-\text{path}})$*. This shows that the serial consensus controller with* $L_1 = L_{\text{ahead}-\text{path}}$ *and* $L_2 = L_{\text{behind}-\text{path}}$ *is* 1*-step implementable and only requires local relative feedback.*

On the other hand, if $L_2L_1 \notin \mathcal{A}^1(W,c)$, then another alternative is to extend the local measurements. This is the case when $L_1 = L_2 = L_{\text{ahead}-\text{path}}$, because then we have

$$L^2_{\text{ahead}-\text{path}} = \begin{bmatrix} 0 & 0 & & & \\ -1 & 1 & & & \\ 1 & -2 & 1 & & \\ & \ddots & \ddots & \ddots & \\ & & 1 & -2 & 1 \end{bmatrix},$$

where $L^2_{\text{ahead}-\text{path}} \notin \mathcal{A}^1(W_{\text{ahead}-\text{path}}, c)$. Then, both in this particular case and in general, it is sufficient to add measurements to neighbors' neighbors. That is, given $L_1, L_2 \in \mathcal{A}^1(W,c)$, then it holds that the product $(L_2L_1) \in \mathcal{A}^2(W,c')$, as guaranteed by Proposition 1. The sum will clearly satisfy $(L_2 + L_1) \in \mathcal{A}^1(W,2c)$. This idea of including extra measurements to improve coordination performance has been used in e.g. [Darbha et al., 2019] and is thus not new to the vehicular formation literature. However, in the conventional consensus, such an addition of a bounded number of neighbors does not provide improved scaling of performance in N in the same way as serial consensus does.

5. Examples

This section will provide three examples to illustrate our main results and how serial consensus compares to the conventional consensus protocol. We choose the same velocity and positional gains for the following examples to make them comparable. However, the illustrated fundamental differences between the protocols are all independent of this choice.

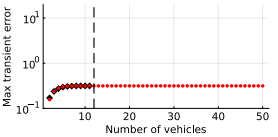

(a) Serial (red) with directed cycle feedback. Conventional (black) with directed cycle feedback. Dashed line marks when conventional becomes unstable.

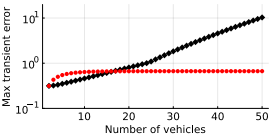

(b) Serial (red) with directed path feedback. Conventional (black) with directed path feedback.

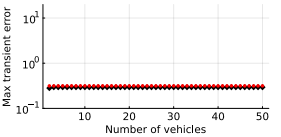

(c) Serial (red) with bidirectional feedback. Conventional (black) with symmetric position and asymmetric velocity feedback.

Figure 3. Maximum transient error is plotted against the number of vehicles. The serial consensus converges to a fixed transient in all cases. The conventional is shown to have three characteristic behaviors: instability, rapidly growing transient, and bounded transient.

5.1 Scalable stability

Example 3

Consider the uni-directional circular graph structure with the graph Laplacian defined as

$$(L_{\text{ahead}-\text{cycle}})_{i,j} = \begin{cases} 1 & \text{if } i = j \\ -1 & \text{if } i = j+1 \pmod N \end{cases}.$$

The conventional consensus protocol is then

$$\ddot{x} = -r_{\text{v}} L_{\text{ahead}-\text{cycle}} \dot{x} - r_{\text{p}} L_{\text{ahead}-\text{cycle}} x + u_{\text{ref}}$$

This system is known to be troublesome since unless r_{p} and r_{v} are chosen to depend on the number of vehicles N, the closed-loop system will be unstable for sufficiently large N. Algebraically, this is a consequence of the smallest in magnitude eigenvalues of $L_{\text{ahead}-\text{cycle}}$ approaching the origin at an angle to the real axis as N grows [Tegling et al., 2023].

On the other hand, the serial consensus protocol

$$\ddot{x} = -(p_1 + p_2) L_{\text{ahead}-\text{cycle}} \dot{x} - p_1 p_2 L^2_{\text{ahead}-\text{cycle}} x + u_{\text{ref}}$$

is stable for any N as long as p_1 and p_2 are chosen to be positive, which follows from Theorem 1. Provided $p_1 \neq p_2$, Theorem 2 further asserts that it has scalable performance with respect to the graph Laplacian $L_{\text{ahead}-\text{cycle}}$.

A comparison of the transient behavior for the two formations is shown in Figure 2. The figure shows how the formation controlled through the serial consensus protocol is stable and has similar behavior independent of the number of vehicles. Meanwhile, the one controlled with conventional consensus performs similarly for small N but eventually loses stability for large N. The maximum transient error for multiple N is shown in 3a. The figure illustrates how the maximum deviations for small N are very similar between the two systems until the conventional consensus eventually becomes unstable.

5.2 Scalable performance

Example 4

Here, we will illustrate the significant difference in performance between conventional and serial consensus in the case of a directed vehicle string. For this purpose, consider the directed path topology whose graph Laplacian is given by (11). In this case, it is easy to verify that the conventional consensus protocol

$$\ddot{x} = -r_{\mathrm{v}} L_{\mathrm{ahead-path}} \dot{x} - r_{\mathrm{p}} L_{\mathrm{ahead-path}} x + u_{\mathrm{ref}}$$

will stabilize the vehicle formation for any choice of positive r_{v} and r_{p}. However, the transient behavior will scale poorly independent of the choice of r_{v} and r_{p} as is illustrated in Figure 2. The formation under this control lacks string stability [Seiler et al., 2004]; disturbances propagate and grow along the string. On the other hand, the serial consensus protocol with

$$\ddot{x} = -(p_1 + p_2) L_{\mathrm{ahead-path}} \dot{x} - p_1 p_2 L^2_{\mathrm{ahead-path}} x + u_{\mathrm{ref}}$$

will have scalable performance with respect to the position error $e_{\mathrm{p}} = d + L_{\mathrm{ahead-path}} x$ and velocity error e_{v} as long as $p_1 \neq p_2$. This is illustrated in Figure 2g and Figure 2h, where the same initial conditions and parameters are used for both consensus protocols. From the figures, it is clear that both formations are stable. However, the conventional consensus protocol has a much larger transient than the serial consensus protocol, which worsens as the number of vehicles increases. The maximum transient error is shown in Figure 3b. The maximum transient grows exponentially in N for the conventional protocol. Meanwhile, the corresponding transient for the serial consensus reaches a ceiling, as is predicted by Theorem 2.

5.3 Different graph Laplacians

Example 5

The conventional consensus design has been shown to have acceptable performance in a vehicle string when different Laplacians are used in the position and velocity feedback [Herman et al., 2017]. In particular, the use of the directed path graph Laplacian for the velocity term (look-ahead) and an undirected path graph Laplacian for the positional term (look-ahead and look-behind), defined as

$$(L_{\mathrm{undir-path}})_{i,j} = \begin{cases} 2 & \text{if } i = j \text{ and } 2 \leq i \leq N-1 \\ 1 & \text{if } i = j \text{ for } i = 1, N \\ -1 & \text{if } |i-j| = 1. \end{cases}$$

The resulting closed-loop system is then

$$\ddot{x} = -r_{\mathrm{v}} L_{\mathrm{ahead-path}} \dot{x} - r_{\mathrm{p}} L_{\mathrm{undir-path}} x + u_{\mathrm{ref}}. \tag{13}$$

The step responses for $N = 10, 100$ can be seen in Figure 4a and 4b.

This can be compared to the serial consensus utilizing bidirectional information. For instance, if the forward-looking graph Laplacian $L_{\text{ahead}-\text{path}}$ from (11) is used together with the corresponding backward-looking graph Laplacian $L_{\text{behind}-\text{path}}$ from (12). The resulting closed-loop system is then

$$\ddot{x} = -(p_2 L_{\text{ahead}-\text{path}} + p_1 L_{\text{behind}-\text{path}})\dot{x} - p_1 p_2 L_{\text{ahead}-\text{path}} L_{\text{behind}-\text{path}} x + u_{\text{ref}}. \quad (14)$$

The step responses are in Figure 4. From the figure, we can observe that the serial and conventional consensus can have similar transient performance. This is also illustrated by the maximum transient error for various N, as shown in Figure 3c. In this case, both the serial and conventional consensus seem to have a maximum transient error that is independent of N. Even if the two protocols can have similar performance, we find it easier to predict that the serial consensus will have good performance than various versions of the conventional protocol. Indeed, the protocol proposed by [Herman et al., 2017] is similar to the serial consensus controller in Example 4.

6. Conclusions and Directions for Future Work

In this paper, we have introduced the serial consensus controller, a distributed formation controller that achieves scalable stability, performance, and robustness. Here, scalability refers to the fact that these properties are independent of the formation size. The performance result is particularly important since it linearly bounds the $\|\cdot\|_\infty$-gain from the initial local errors and reference velocity deviations to the transient local errors and reference velocity deviations, measured in the same norm. This quantity, rather than, for example, an L_2 gain, is directly related to the control and performance objectives. It is also worth noting that these results are achieved with only local relative measurements and linear feedback. The results are particularly interesting for large vehicle platoons where short inter-vehicle distances are desired, and the transient behavior of the platoon is of great importance, though there are strict topological constraints (typically, those of a directed string). However, by the generality of the presented results, they could also be of interest to other networked systems, such as power grids, sensor networks, or multi-robot networks. Since the serial consensus design can guarantee both consensus and bounded transients, it can be seen as a suitable parametrization for optimal control where a potentially different objective is considered.

There are several interesting directions for future work. First, since the serial consensus may require an additional communication step, an open question is whether this can be avoided, perhaps by using local estimators. A second direction is to further investigate the robustness of serial consensus, particularly concerning the term $L^2 x$, which may experience time delays when implemented through message passing. The handling of other disturbances is also of interest. Preliminary results

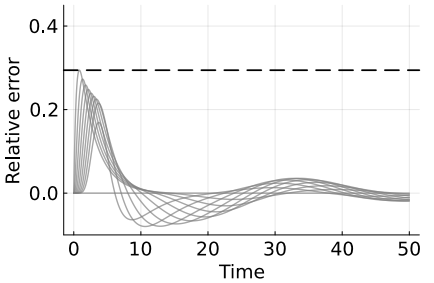

(a) Conventional consensus with symmetric position, asymmetric velocity feedback and $N = 10$ and transient $\sup_t \|e_{\mathrm{p}}\|_\infty = 0.294$.

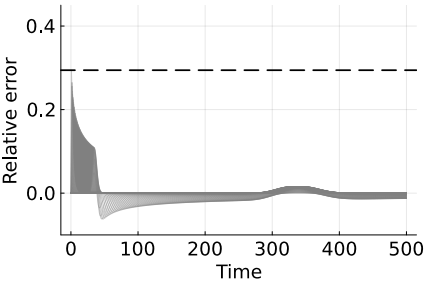

(b) Conventional consensus with symmetric position, asymmetric velocity feedback and $N = 100$ and transient $\sup_t \|e_{\mathrm{p}}\|_\infty = 0.294$.

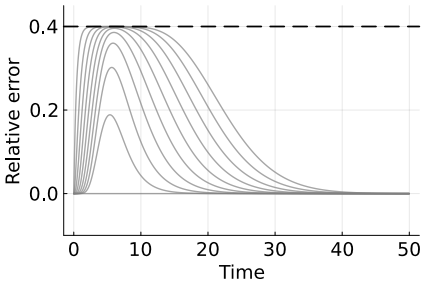

(c) Serial consensus with bidirectional feedback and $N = 10$ and transient $\sup_t \|e_{\mathrm{p}}\|_\infty = 0.400$.

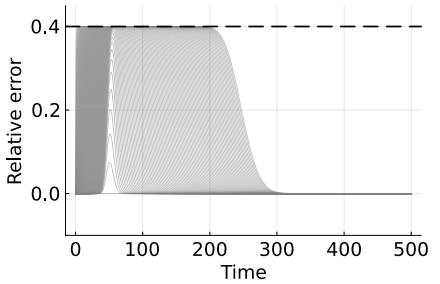

(d) Serial consensus with bidirectional feedback and $N = 100$ and transient $\sup_t \|e_{\mathrm{p}}\|_\infty = 0.400$.

Figure 4. Simulation of the initial value response to $x = 0$, $\dot{x}_{i\neq 1}(0) = 0$, and $\dot{x}_1(0) = 1$. The conventional consensus system (13) is considered with $r_{\mathrm{v}} = 2.5$ and $r_{\mathrm{p}} = 1$, while for the serial consensus system (14), $p_1 = 2$ and $p_2 = 0.5$ are used. The results illustrate that for some graph Laplacians, the serial and conventional consensus can have comparable performance.

[Hansson and Tegling, 2024] have shown that serial consensus can reject certain load disturbances while preserving string stability. Extending the class of permissible disturbances remains an important area for future research. Third, the ability of vehicles to enter and leave a formation necessitates the study of time-varying graph topologies. Finally, implementing serial consensus on a physical system is an exciting next step.

Appendix

A The case $p_1 = p_2$

Suppose $p_1 = p_2 = p$ in (10); then we are interested in the system

$$\ddot{x} = -2pL\dot{x} - p^2L^2x + u_{\text{ref}}.$$

Since the performance result should apply to an arbitrary L, we may consider $\hat{L} = L/p$. Thus, any bound should also be true for

$$\ddot{x} = -2p\hat{L}\dot{x} - p^2\hat{L}^2x + u_{\text{ref}} = -2L\dot{x} - L^2x + u_{\text{ref}}$$

and $p = 1$ can be assumed WLOG. If $L = L_{\text{behind}-\text{path}}$ is used together with the initial condition $e_{\text{v}}(0) = [\mathbf{1}_{N-1}^T, 0]^T$ and $e_{\text{p}}(0) = \mathbf{0}$, the solution will be

$$e_{\text{p}}(t) = tL_{\text{behind}-\text{path}}e^{-tL_{\text{behind}-\text{path}}}e_{\text{v}}(0).$$

The first $N-1$ errors $\tilde{e}_p$ are then

$$\tilde{e}_{\text{p}} = -tJ_{-1}e^{J_{-1}t}\mathbf{1},$$

where $J_{-1} \in \mathbb{R}^{N-1\times N-1}$ is a Jordan block with eigenvalue -1. The matrix exponential is known for the Jordan block, and the exact solution is

$$\tilde{e}_p = te^{-t}\begin{bmatrix} t^{N-2}/(N-2)! \\ t^{N-3}/(N-3)! \\ \vdots \\ 1/0! \end{bmatrix}$$

Using Stirling's Approximation formula to upper bound the factorial, using $t = N-2$, together with some simple bounds, show that $\|e_{\text{p}}\|_\infty \geq \sqrt{N-1}/(4e) \geq \|e_{\text{v}}(0)\|_\infty\sqrt{N-1}/(4e)$. Since the lower bound grows unboundedly, no upper bound can be independent of the network size N. Therefore, we cannot give scalable performance guarantees in the case $p_1 = p_2 = p$ and must require these constants to be distinct.

B Proof of Theorem 1

Proof. i) Any square matrix can be unitarily transformed to upper triangular form by the Schur triangularization theorem. Let $U_kL_kU_k^H = T_k$ be upper triangular. Then the block diagonal matrix $U = \text{diag}(U_1, U_2)$ is a unitary matrix that upper triangularizes A in (8). For any triangular matrix, the eigenvalues lie on the diagonal, which will be the eigenvalues of each $-L_k$. The result follows.

ii) First, consider the closed-loop dynamics of (6)

$$X(s) = (sI + L_1)^{-1}(sI + L_2)^{-1}U_{\text{ref}}(s).$$

Since, U_{ref} is stable, we know that the limit $\lim_{s\to 0} U_{\text{ref}}(s) = U_{\text{ref}}(0)$ exists. To prove that the system achieves second-order consensus, we want to show that

$$\lim_{t\to\infty} y(t) = \lim_{s\to 0} C(s)X(s) = 0$$

for some transfer matrix $C(s)$, which encodes the consensus states. But since the reference dependence is only related to $U_{\text{ref}}(0)$, we can simplify the problem to only consider impulse responses, which has the same transfer function as the initial value response where $\xi_2(0) = U_{\text{ref}}(0)$. Therefore, WLOG, assume that $U_{\text{ref}}(s) = 0$ and an arbitrary initial condition $\xi(0) = [\xi_1(0)^T, \xi_2^T(0)]^T$. The solution of (8) is given by $\exp(At)\xi(0) = S\exp(J(A)t)S^{-1}\xi(0)$ where $J(A)$ is the Jordan normal form of A and S is an invertible matrix. From i) and the diagonal dominance of the graph Laplacians, we know that all eigenvalues of A lie in the left half plane. By Assumption 1, it follows that the zero eigenvalue for each L_k is simple. Now, we prove that these two zero eigenvalues correspond to a Jordan block of size 2. Let $\mathbf{e}_1 = \begin{bmatrix}\mathbf{1}^T & 0\end{bmatrix}^T$ and $\mathbf{e}_2 = \begin{bmatrix}0 & \mathbf{1}^T\end{bmatrix}^T$. Then e_1 is an eigenvector since $A\mathbf{e}_1 = 0$. Now, since $e_1 = Ae_2$ combined with e_1 and e_2 being linearly independent, it follows that they form a Jordan block of size 2 with an invariant subspace spanned by the vectors e_1 and e_2. All other eigenvectors make up an asymptotically stable invariant subspace, it follows that $\xi(t)$ will converge towards a solution in $\text{span}(e_1, e_2)$ and thus $\lim_{t\to\infty} \xi_k(t) = \alpha_k(t)\mathbf{1}$. From $x(t) = \xi_1(t)$ we get $\lim_{t\to\infty} x(t) = \alpha_1(t)\mathbf{1}$, and furthermore, since

$$\dot{x} = \dot{\xi}_1 = -L_1\xi_1 + \xi_2 \to \xi_2 \text{ as } t \to \infty,$$

it follows that $\lim_{t\to\infty} \dot{x}(t) = \alpha_2(t)\mathbf{1}$, which shows that the system achieves second-order consensus. □

C Proof of Proposition 2.

Proof. To prove this we must show that both $A_{\text{pos}}, A_{\text{vel}} \in \mathcal{A}^2(W, c')$ where $A_{\text{pos}} = -L_2L_1$ and $A_{\text{vel}} = -(L_2 + L_1)$. First, A_{pos} and A_{vel} are shown to represent relative feedback. Since $L_1, L_2 \in \mathcal{A}^1(W, c)$, it holds that $A_{\text{pos}}\mathbf{1} = -L_2L_1\mathbf{1} = -L_20 = 0$ and similarly $A_{\text{vel}}\mathbf{1} = -(L_2 + L_1)\mathbf{1} = 0$.

Second, we show that the gain is bounded. For the positional feedback, we have

$$\|A_{\text{pos}}\|_\infty = \| - L_2L_1\|_\infty \leq \|L_2\|_\infty\|L_1\|_\infty \leq c^2,$$

which followed from the submultiplicativity of the induced norm. For the velocity feedback

$$\|A_{\text{vel}}\|_\infty = \| - (L_2 + L_1)\|_\infty \leq \|L_2\|_\infty + \|L_1\|_\infty \leq 2c,$$

where the triangle inequality was utilized. Let $c' = \max\{2c, c^2\}$, then clearly it holds true that both $\|A_{\text{pos}}\|_\infty, \|A_{\text{vel}}\|_\infty \leq c'$.

Finally, we consider the sparsity pattern. Since W is non-negative, W^2 is also non-negative. Next, since adding non-negative elements to a matrix cannot remove any positive elements, we get the following implications

$$\left[I+W+W^2\right]_{i,j}=0 \Longrightarrow [I+W]_{i,j}=0 \Longrightarrow [-L_1-L_2]_{i,j}=0,$$

which follows from the definition of $L_1, L_2 \in \mathcal{A}^1(W,c)$.

To show that $A_{\text{pos}} = -L_2L_1$ will be sparse, we note that any graph Laplacian can be written as $L = D - W$, where D is a diagonal matrix and W is non-negative. Thus, we will consider the sparsity of

$$-(D_2 - W_2)(D_1 - W_1) = -D_2D_1 + D_2W_1 + W_2D_1 - W_2W_1.$$

Since multiplication with a diagonal matrix preserves sparsity, it holds that the first three terms satisfy

$$[I+W]_{i,j}=0 \Longrightarrow [-D_2D_1 + D_2W_1 + W_2D_1]_{i,j}=0.$$

For W_2W_1 we can introduce a non-negative matrix E such that $[W_{1/2}+E]_{i,j} = 0 \iff W_{i,j} = 0$ and in particular this construction ensures that $[(W_1+E)(W_2+E)]_{i,j} = 0 \iff [W^2]_{i,j} = 0$. Consider the following expanded product

$$(W_2+E)(W_1+E) = W_2W_1 + W_2E + EW_1 + E^2,$$

where all terms are products of non-negative matrices and are therefore also non-negative. It thus holds that $[W^2]_{i,j} \implies [W_2W_1]_{i,j} = 0$. Combining the results for A_{pos} shows that

$$\left[I+W+W^2\right]_{i,j}=0 \implies [L_2L_1]_{i,j} = 0.$$

This concludes the proof. □

Acknowledgements

The authors are with the Department of Automatic Control and the ELLIIT Strategic Research Area at Lund University, Lund, Sweden. This work was partially funded by Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation and the Swedish Research Council through Grant 2019-00691.

References

Bamieh, B., M. R. Jovanovic, P. Mitra, and S. Patterson (2008). "Effect of topological dimension on rigidity of vehicle formations: fundamental limitations of local feedback". In: *Proc. IEEE Conf. Decis. Control*, pp. 369–374.

Bamieh, B., M. R. Jovanovic, P. Mitra, and S. Patterson (2012). "Coherence in large-scale networks: dimension-dependent limitations of local feedback". *IEEE Trans. Autom. Control* **57**:9, pp. 2235–2249.

Barooah, P., P. G. Mehta, and J. P. Hespanha (2009). "Mistuning-based control design to improve closed-loop stability margin of vehicular platoons". *IEEE Trans. Autom. Control* **54**:9, pp. 2100–2113.

Besselink, B. and S. Knorn (2018). "Scalable input-to-state stability for performance analysis of large-scale networks". *IEEE Control Syst. Lett.* **2**:3, pp. 507–512.

Cantos, C., J. Veerman, and D. Hammond (2016). "Signal velocity in oscillator arrays". *Eur. Phys. J. Spec.* **225**:6, pp. 1115–1126.

Darbha, S., S. Konduri, and P. R. Pagilla (2019). "Benefits of v2v communication for autonomous and connected vehicles". *IEEE Trans. Intell. Transp. Syst.* **20**:5, pp. 1954–1963.

Dolk, V. S., J. Ploeg, and W. P. M. H. Heemels (2017). "Event-triggered control for string-stable vehicle platooning". *IEEE Trans. Intell. Transp. Syst.* **18**:12, pp. 3486–3500.

Fax, J. A. and R. M. Murray (2004). "Information flow and cooperative control of vehicle formations". *IEEE Trans. Autom. Control* **49**:9, pp. 1465–1476.

Feintuch, A. and B. Francis (2012). "Infinite chains of kinematic points". *Automatica* **48**:5, pp. 901–908.

Feng, S., Y. Zhang, S. E. Li, Z. Cao, H. X. Liu, and L. Li (2019). "String stability for vehicular platoon control: definitions and analysis methods". *Annu. Rev. Control* **47**, pp. 81–97.

Gunter, G., D. Gloudemans, R. E. Stern, S. McQuade, R. Bhadani, M. Bunting, M. L. Delle Monache, R. Lysecky, B. Seibold, J. Sprinkle, B. Piccoli, and D. B. Work (2021). "Are commercially implemented adaptive cruise control systems string stable?" *IEEE Trans. Intell. Transp. Syst.* **22**:11, pp. 6992–7003.

Hansson, J. and E. Tegling (2023). "A closed-loop design for scalable high-order consensus". In: *Proc. IEEE Conf. Decis. Control*, pp. 7388–7394.

Hansson, J. and E. Tegling (2024). "Performance bounds for multi-vehicle networks with local integrators". *arXiv:2410.14525*.

Hao, H. and P. Barooah (2012). "On achieving size-independent stability margin of vehicular lattice formations with distributed control". *IEEE Trans. Autom. Control* **57**:10, pp. 2688–2694.

Herman, I., S. Knorn, and A. Ahlén (2017). "Disturbance scaling in bidirectional vehicle platoons with different asymmetry in position and velocity coupling". *Automatica* **82**, pp. 13–20.

Jadbabaie, A., J. Lin, and A. S. Morse (2003). "Coordination of groups of mobile autonomous agents using nearest neighbor rules". *IEEE Trans. Autom. Control* **48**:6, pp. 988–1001.

Oh, K.-K., M.-C. Park, and H.-S. Ahn (2015). "A survey of multi-agent formation control". *Automatica* **53**, pp. 424–440.

Olfati-Saber, R. and R. M. Murray (2004). "Consensus problems in networks of agents with switching topology and time-delays". *IEEE Trans. Autom. Control* **49**:9, pp. 1520–1533.

Pasqualetti, F., S. Zampieri, and F. Bullo (2014). "Controllability metrics, limitations and algorithms for complex networks". In: *Proc. Am. Control Conf.,* pp. 3287–3292.

Patterson, S. and B. Bamieh (2014). "Consensus and coherence in fractal networks". *IEEE Trans. Control Netw. Syst.* **1**:4, pp. 338–348.

Ren, W. and E. Atkins (2007). "Distributed multi-vehicle coordinated control via local information exchange". *Int. J. Robust Nonlinear Control* **17**:10-11, pp. 1002–1033.

Seiler, P., A. Pant, and K. Hedrick (2004). "Disturbance propagation in vehicle strings". *IEEE Trans. Autom. Control* **49**:10, pp. 1835–1842.

Siami, M. and N. Motee (2016). "Fundamental limits and tradeoffs on disturbance propagation in large-scale dynamical networks". *IEEE Trans. Autom. Control* **61**:12, pp. 4055–4062.

Stüdli, S., M. Seron, and R. Middleton (2017a). "From vehicular platoons to general networked systems: string stability and related concepts". *Annu. Rev. Control* **44**, pp. 157–172.

Stüdli, S., M. M. Seron, and R. H. Middleton (2017b). "Vehicular platoons in cyclic interconnections with constant inter-vehicle spacing". *IFAC-PapersOnLine* **50**:1, pp. 2511–2516.

Swaroop, D. and J. Hedrick (1996). "String stability of interconnected systems". *IEEE Trans. Autom. Control* **41**:3, pp. 349–357.

Tegling, E., B. Bamieh, and H. Sandberg (2023). "Scale fragilities in localized consensus dynamics". *Automatica* **153**, p. 111046.

Tegling, E., P. Mitra, H. Sandberg, and B. Bamieh (2019). "On fundamental limitations of dynamic feedback control in regular large-scale networks". *IEEE Trans. Autom. Control* **64**:12, pp. 4936–4951.

Wijnbergen, P., M. Jeeninga, and B. Besselink (2021). "Nonlinear spacing policies for vehicle platoons: a geometric approach to decentralized control". *Syst. Control Lett.* **153**, p. 104954.

Zhou, K. and J. C. Doyle (1998). *Essentials of robust control*. Vol. 104. Prentice hall Upper Saddle River, NJ.

Paper III

Performance Bounds for Multi-Vehicle Networks with Local Integrators

Jonas Hansson **Emma Tegling**

Abstract

In this work, we consider the problem of coordinating a collection of n^{th}-order integrator systems. The coordination is achieved through the novel serial consensus design; this control design achieves a stable closed-loop system while adhering to the constraint of only using local and relative measurements. Earlier work has shown that second-order serial consensus can stabilize a collection of double integrators with scalable performance conditions independent of the number of agents and topology. This paper generalizes these performance results to an arbitrary order $n \geq 1$. The derived performance bounds depend on the condition number, measured in the vector-induced maximum matrix norm, of a general diagonalizing matrix. We precisely characterize how a minimal condition number can be achieved. Third-order serial consensus is illustrated through a case study of PI-controlled vehicular formation, where the added integrators are used to mitigate the effect of unmeasured load disturbances. The theoretical results are illustrated through examples.

1. Introduction

Control of complex systems is a field with a rich literature. A few notable examples are control of energy networks, multi-agent coordination, and transportation networks. The common theme is that rich behavior emerges from the many interconnections between subsystems. This work concerns the problem of multi-agent coordination, a problem pioneered by [Fax and Murray, 2004; Olfati-Saber and Murray, 2004; Jadbabaie et al., 2003], and the sub-problem of distributed consensus. The consensus protocol has seen many significant theoretical contributions and generalizations. For instance, the works [Bamieh et al., 2012; Patterson and Bamieh, 2014; Siami and Motee, 2014; Pates et al., 2017] have considered the sensitivity to noise, often studied through the so-called coherence. There are many more results, such as resilience, scale-fragility, and scalable performance, to name a few.

The consensus protocol has also been generalized to achieve more complex coordination. In [Ren and Atkins, 2005], the second-order consensus protocol was introduced. This protocol simultaneously coordinates velocity and position for a network of double integrators. In its simplest form, it can be formulated as $\ddot{x} = -r_{\text{vel}}L\dot{x} - r_{\text{pos}}Lx + u_{\text{ref}}$, where $r_{\text{vel}}, r_{\text{pos}}$ are positive constants and L is the graph Laplacian. Since the seminal work [Ren and Atkins, 2005], the protocol has been extensively studied; for example, in terms of coherence [Patterson and Bamieh, 2014]; stability [Stüdli et al., 2017b]; string stability [Stüdli et al., 2017a; Feng et al., 2019].

Following the idea of second-order consensus, a corresponding high-order version was considered in [Ren et al., 2007]:

$$x^{(n)} = -r_{(n-1)}Lx^{(n-1)} - \cdots - r_{(1)}L\dot{x} - r_{(0)}Lx + u_{\text{ref}},$$

where $r_{(1)} = r_{\text{vel}}$, $r_{(0)} = r_{\text{pos}}$. Like the second-order protocol, this model has received some attention. An interesting feature of the high-order consensus protocol is that it lacks *scalable stability* [Tegling et al., 2023]; that is, the weights $r_{(k)}$ need to be tuned for the specific graph Laplacian that is used to avoid instability when the number of agents grows.

A more recent variation of the high-order consensus protocol is *serial consensus*, which was proposed in [Hansson and Tegling, 2023]. In the Laplace domain, the serial consensus system is $(\prod_{k=1}^{n} sI + p_k L)X(s) = U_{\text{ref}}(s)$. In the second-order case, the corresponding time-domain representation is $\ddot{x} = -(p_1 + p_2)L\dot{x} - p_2 p_1 L^2 x + u_{\text{ref}}$. The idea is to design a stable closed-loop system independently of the graph Laplacian to mitigate the lack of scalable stability in the conventional high-order consensus [Stüdli et al., 2017b]. However, in high-order consensus, stability is not the only concern, but to be practical, it is often necessary that the closed-loop system satisfies additional performance criteria, a problem also considered in [Macellari et al., 2017]. Here, we are in particular concerned with *scalable performance* criteria, meaning that the transient should not grow with the size of the formation. This prevents problems such as string instability, a well-known problem

in the vehicle formation literature; two surveys are [Stüdli et al., 2017a; Feng et al., 2019]. In our work [Hansson and Tegling, 2024], we showed that second-order serial consensus can be used to achieve scalable performance and thus avoid string instability, regardless of network topology. These benefits come at the cost of at most one additional 'hop' of communication (in general, at most an n-hop neighborhood is required).

Here, we continue our work on the serial consensus system. An outstanding question was whether the performance result generalizes to higher-order consensus protocols in general and the practically relevant third-order protocol in particular. The main contribution of this work is a positive answer to this question: high-order serial consensus achieves scalable performance independent of topology. This significantly improves over conventional high-order consensus, which does not even exhibit scalable stability.

Although second-order consensus systems are partially motivated by the vehicle formation problem and multi-agent coordination, they cannot reject unmeasured load disturbances, which unavoidably occur in a real-world setting. Motivated by the classic solution of introducing integral control to counteract steady-state errors, we illustrate how to introduce integral action through a third-order serial consensus design. We show that this system can mitigate load disturbances with scalable performance in terms of maximum transient state deviations due to arbitrary initial conditions. Our scalable performance result is quantified through an exact and novel characterization of the minimal scaled condition number of a diagonalizing matrix: $\inf_{A=SDS^{-1}} \|S\|_\infty \|S^{-1}\|_\infty$.

The article is structured as follows. Section 2 will introduce the mathematical notation and models used throughout the article. Then, our main theorem and lemma are presented in Section 3. The results are later adapted in Section 4 to the setting of PI-controlled vehicle formations to show the scalable performance thereof. The vehicle formation results are further elaborated through numerical simulations in Section 5. Finally, we conclude the paper with our conclusions and some discussion in Section 6.

2. Preliminaries

2.1 Definitions and network model

We will by $\mathcal{G}(\mathcal{V},\mathcal{E})$ denote a graph with $N=|\mathcal{V}|$ vertices and with edge set $\mathcal{E}\subset\mathcal{V}\times\mathcal{V}$. The graph can also be represented by the weighted adjacency matrix W, which satisfies $W_{i,j}>0 \iff (j,i)\in\mathcal{V}$. The graph is called undirected if $W^\top = W$. The graph Laplacian L is defined as $L=D-W$ where $D=\mathrm{diag}(W\mathbf{1})$.

A graph contains a directed spanning tree if, for some node i, a path exists to all other vertices $j\in\mathcal{V}\setminus i$. If the graph contains a directed spanning tree, then the associated graph Laplacian has a unique zero eigenvalue, and the rest of the eigenvalues have strictly positive real parts.

Throughout this paper we will denote the standard vector ∞-norm of $x \in \mathbb{C}^N$ as $\|x\|_\infty = \max_i |x_i|$. We will reuse the notation for the corresponding induced matrix norm of $C \in \mathbb{C}^{M\times N}$, that is, $\|C\|_\infty = \sup_{\|x\|_\infty=1} \|Cx\|_\infty = \max_i \sum_{j=1}^N |C_{ij}|$. We denote the standard Kronecker product between two matrices $A \otimes B$.

2.2 Serial Consensus

We begin by introducing serial consensus.

Definition 1—n^{TH}-order serial consensus
The following closed-loop system, expressed in the Laplace domain,

$$\left(\prod_{k=1}^{n} sI + L_k\right) X(s) = U_{\text{ref}}(s) \tag{1}$$

is called the n^{th}-order serial consensus system.

This can be compared to the conventional n^{th}-order consensus system which in the case $n = 2$ is

$$\ddot{x} = -L_{\text{vel}}\dot{x}(t) - L_{\text{pos}}x(t) + u_{\text{ref}}(t).$$

The corresponding second-order serial consensus system is

$$\ddot{x}(t) = -(L_2 + L_1)\dot{x} - L_2L_1x + u_{\text{ref}}.$$

The two protocols are similar and can be seen as two natural ways of generalizing the regular continuous-time consensus protocol, that is, $\dot{x} = -Lx$, preserving the limitation of only utilizing relative feedback. We refer to [Hansson and Tegling, 2023] for a more elaborate discussion.

One practical property of the serial consensus is that there exists a simple sufficient condition for stability; under the assumption that each of the graph Laplacians L_k contain a directed spanning tree, the serial consensus will eventually achieve coordination [Hansson and Tegling, 2023]. When identical graph Laplacians are used, i.e., $L_k = p_kL$, additional performance results can be derived. This setting is what we will study in this work. In this case, the serial consensus system can be represented in the following state-space form

$$\begin{bmatrix} \dot{x} \\ \ddot{x} \\ \vdots \\ x^{(n)} \end{bmatrix} = \begin{bmatrix} & I & & \\ & & \ddots & \\ & & & I \\ -a_0L^n & -a_1L^{n-1} & \dots & -a_{n-1}L \end{bmatrix} \begin{bmatrix} x \\ \dot{x} \\ \vdots \\ x^{(n-1)} \end{bmatrix} + \begin{bmatrix} \\ \\ \\ u_{\text{ref}} \end{bmatrix},$$

where a_k are identified with the weights p_k through the relation $\prod_{k=1}^n (s+p_k) = \sum_{k=0}^n a_k s^k$. To study the evolution of relative errors, we will consider a transformed

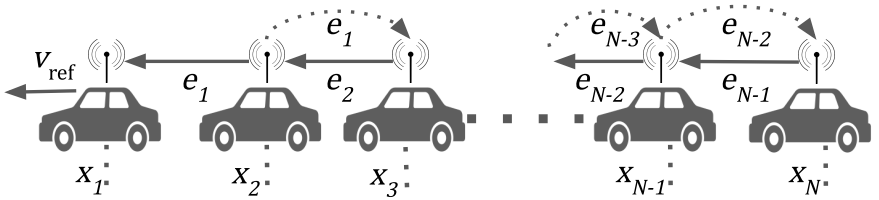

Figure 1. Illustration of vehicle formation control through serial consensus. In this example of a directed string formation, each vehicle measures the relative distance to its neighbor, and message passes its measurements to its follower.

state-space representation that is particularly useful for this purpose

$$\underbrace{\begin{bmatrix} L^{n-1}\dot{x} \\ L^{n-2}\ddot{x} \\ \vdots \\ x^{(n)} \end{bmatrix}}_{\dot{\xi}} = \underbrace{\begin{bmatrix} & L & & \\ & & \ddots & \\ & & & L \\ -a_0 L & -a_1 L & \dots & -a_{n-1} L \end{bmatrix}}_{\mathcal{A}=A\otimes L} \underbrace{\begin{bmatrix} L^{n-1}x \\ L^{n-2}\dot{x} \\ \vdots \\ x^{(n-1)} \end{bmatrix}}_{\xi} + \begin{bmatrix} \\ \\ \\ u_{\text{ref}} \end{bmatrix}, \tag{2}$$

where

$$A = \begin{bmatrix} & 1 & & \\ & & \ddots & \\ & & & 1 \\ -a_0 & -a_1 & \dots & -a_{n-1} \end{bmatrix}. \tag{3}$$

Here, we note that A is related to the controllable canonical form associated with the transfer function $\prod_{k=1}^{n}(s+p_k)^{-1}$.

Unlike conventional consensus, implementing serial consensus requires information on the state of neighbors' neighbors. If each Laplacian L_k of (1) represents locally available information, each agent must access data from their n-hop neighborhood. This can be achieved through either local message passing or by constructing L and the communication so that each agent's required relative measurements are locally available. The message-passing procedure is illustrated in Figure 1. See [Hansson and Tegling, 2023] for further details.

2.3 Scalable performance

The performance of a system can be evaluated in multiple ways. Inspired by the vehicle formation setting where velocity tracking and relative distance keeping is required, we define the measurements $e_{\text{p}} = L(x-d)$ and $e_{\text{v}} = \dot{x} - v_{\text{ref}}\mathbf{1}$ with $d \in \mathbb{R}^n$ and $v_{\text{ref}} \in \mathbb{R}$ being the quantities that should be kept small throughout the dynamic phase and in the presence of bounded disturbances $w(t)$. This leads to the following.

Definition 2—Scalable Performance
A formation controller defined for any formation size N that ensures

$$\sup_{t\geq 0}\left\|\begin{bmatrix}e_{\mathrm{p}}(t)\\ e_{\mathrm{v}}(t)\end{bmatrix}\right\|_{\infty}\leq \alpha_e\left\|\begin{bmatrix}e_{\mathrm{p}}(0)\\ e_{\mathrm{v}}(0)\end{bmatrix}\right\|_{\infty}+\alpha_w\sup_{t\geq 0}\|w(t)\|_{\infty}, \tag{4}$$

where α_e and α_w are constants, independent of the number of agents N, is said to achieve scalable performance.

The choice of the norm here is important, as the maximum deviations from equilibrium determine whether collisions will occur or speed limits will be violated. We note that this definition is closely related to $\mathcal{L}_\infty$ string stability as described in [Ploeg et al., 2014], with the main difference being that our definition considers amplification from initial measurements rather than initial absolute states. This distinction makes our definition stricter, as small absolute deviations imply small relative deviations, but the converse is not generally true. A key strength of this definition is that the upper bound depends only on locally available information.

The same definition cannot be directly applied to a more general system since what constitutes good performance is application-dependent. Instead, we consider a more general definition for the high-order consensus protocols.

Definition 3—Scalable Output Performance
A formation controller for a system with outputs y and disturbances w defined for any formation size N that ensures

$$\sup_{t\geq 0}\|y(t)\|_{\infty}\leq \alpha_y\|y(0)\|_{\infty}+\alpha_w\sup_{t\geq 0}\|w(t)\|_{\infty},$$

where α_y and α_w are constants, independent of the number of agents N, is said to achieve scalable output performance.

This definition generalizes that the transient bound should depend only on locally available outputs. Note that w_i can model the effect of additional initial conditions not captured by the output initial conditions. Like scalable performance, this definition is closely related to $\mathcal{L}_\infty$ string stability; in this work, we consider $y=[(L^{n-1}x)^\top,(L^{n-2}\dot{x})^\top,\ldots,(x^{(n-1)})^\top]^\top$, where $\mathcal{L}_\infty$ string stability can be inferred from scalable output stability since $\|y(0)\|_\infty\leq\max\{1,\|L\|^{n-1}\}\|x(0)\|_\infty$. Another related concept is scalable input-to-state stability (sISS) [Besselink and Knorn, 2018]. The main differences are 1) that we consider the evolution of the outputs and 2) that we omit the requirement for a convergence rate independent of the number of agents. In the case of vehicle formations, and using the measurements e_{p} and e_{v}, scalable output performance implies that no vehicle will deviate far from the desired formation and desired velocity at any point in time. We argue these are the main objectives in this setting, while a uniform convergence rate is secondary and potentially not feasible due to locality and gain constraints.

3. Main Results

Our first result shows that the serial consensus will have a bounded transient as a response to an arbitrary initial condition and thus achieves scalable output performance in the absence of external disturbances.

Theorem 1
The n^{th}-order serial consensus system

$$\left(\prod_{k=1}^{n} sI + p_k L\right) X(s) = U(s), \tag{5}$$

with $U(s) = 0$, states $\xi(t) = \left[(L^{n-1}x)^\top, (L^{n-2}\dot{x})^\top, \ldots, (x^{(n-1)})^\top\right]^\top$, and distinct $p_k > 0$ satisfies

$$\sup_{t\geq 0} \|\xi(t)\|_\infty \leq \|S\|_\infty \|S^{-1}\|_\infty \|\xi(0)\|_\infty,$$

where S is any invertible matrix that diagonlizes A in (3).

Proof. A state-space representation of the problem, using the states ξ, is given by (2). The solution to the initial value problem is provided by $\xi(t) = e^{A\otimes Lt}\xi(0)$. Since all p_k are distinct, and that A is associated with the controllable canonical form of the transfer function $\prod_{k=1}^{n}(s+p_k)^{-1}$, it follows that A has n unique negative eigenvalues, being all $-p_k$. Therefore, a diagonalizing and invertible matrix S exists, such that $A = -SPS^{-1}$. The transient can be bounded through the following sequence:

$$\begin{aligned}
\|\xi(t)\|_\infty &= \|e^{A\otimes Lt}\xi(0)\|_\infty \\
&= \|(S\otimes I)e^{-P\otimes Lt}(S^{-1}\otimes I)\xi(0)\|_\infty \\
&\leq \|S\|_\infty \|S^{-1}\|_\infty \|\xi(0)\|_\infty.
\end{aligned}$$

Here, we used submultiplicativity and that $P\otimes L$ is a graph Laplacian, thus satisfying $\|e^{-P\otimes Lt}\|_\infty \leq 1$ for all $t \geq 0$. □

This result shows that the serial consensus system achieves scalable output performance in terms of the outputs $y = \xi$, independent of the graph structure. This is true since the bound is independent of the number of agents N and only depends on the n design parameters p_k. Since the bound holds for any diagonalizing matrix S, there is potential for improvement. The following general lemma identifies the optimal S, which generates the tightest upper bound.

Lemma 1
If $M \in \mathbb{R}^{n\times n}$ has n distinct eigenvalues, then

$$\inf_{M=SDS^{-1}} \|CS\|_\infty \|S^{-1}B\|_\infty = \|CS_*K\|_\infty, \tag{6}$$

where K and D are diagonal matrices, C and B compatible matrices, S_* any matrix satisfying $M = S_*DS_*^{-1}$, and K given by $K_{i,i} = \sum_j |S_*^{-1}B|_{i,j}$.

Proof. We note that if either $C = 0$ or $B = 0$, then the minimum is 0, so from now on, we assume both to be nonzero. Since the eigenvalues of M are distinct, the eigenvectors are uniquely determined up to scaling and permutation. Thus, given one diagonalizing matrix S_*, all others can be obtained through $S = S_*KQ$, where K is invertible and diagonal and Q is a permutation matrix. The minimization problem can be simplified through

$$\begin{aligned}\inf_{M=SDS^{-1}} &\|CS\|_\infty\|S^{-1}B\|_\infty \\ &= \inf_{K,Q}\|CS_*KQ\|_\infty\|Q^\top K^{-1}S_*^{-1}B\|_\infty \\ &= \inf_{K}\|CS_*K\|_\infty\|K^{-1}S_*^{-1}B\|_\infty.\end{aligned}$$

The choice of Q does not impact the size. For K, we note that scaling by any constant $a \neq 0$ preserves the magnitude:

$$\|CS_*aK\|_\infty\|(aK)^{-1}S_*^{-1}B\|_\infty = \|CS_*K\|_\infty\|K^{-1}S_*^{-1}B\|_\infty.$$

Thus, $\|K^{-1}S_*^{-1}B\|_\infty = 1$ can be assumed without loss of generality. Now, let $\hat{K}$ be an arbitrary positive diagonal matrix satisfying $\|\hat{K}^{-1}S_*^{-1}B\|_\infty = 1$. By the definition of the norm, we know that $|\hat{K}_{j,j}^{-1}|\sum_{k=1}^n |S_*^{-1}B|_{j,k} \leq 1$. Another possible choice is $\tilde{K}_{j,j}(\varepsilon) = \sum_{k=1}^n |S_*^{-1}B|_{j,k} + \varepsilon$, where $\varepsilon > 0$. For this choice, $\|[\tilde{K}(\varepsilon)]^{-1}S_*^{-1}B\|_\infty \leq \|\hat{K}^{-1}S_*^{-1}B\|_\infty$ and $|\tilde{K}_{i,i}(\varepsilon)| \leq |\hat{K}_{ii}| + \varepsilon$ for all i. This implies the following

$$\begin{aligned}\|CS_*\tilde{K}(\varepsilon)\|_\infty &= \max_i \sum_{j=1}^n |[CS_*]_{i,j}||\tilde{K}_{j,j}(\varepsilon)| \\ &\leq \max_i \sum_{j=1} |[CS_*]_{i,j}||\hat{K}_{j,j} + \varepsilon| \leq \|CS_*\hat{K}\|_\infty + \varepsilon\|CS_*\|_\infty.\end{aligned}$$

Now, taking the limit

$$\begin{aligned}\|CS_*\tilde{K}(0)\|_\infty &= \lim_{\varepsilon\to 0^+} \|CS_*\tilde{K}(\varepsilon)\|_\infty\|\tilde{K}^{-1}(\varepsilon)S_*^{-1}B)\|_\infty \\ &\leq \lim_{\varepsilon\to 0^+} \|CS_*(\hat{K} + I\varepsilon)\|_\infty = \|CS_*\hat{K}\|_\infty\|\hat{K}^{-1}S_*^{-1}B\|_\infty,\end{aligned}$$

where we used that $B \neq 0$ and that S_*^{-1} has full rank to conclude that $\lim_{\varepsilon\to 0^+} \|\tilde{K}^{-1}(\varepsilon)S_*^{-1}B\|_\infty = 1$. This shows that 1) $\|CS\|_\infty\|S^{-1}B\|_\infty \geq \|CS_*\tilde{K}(0)\|_\infty$ for any diagonalizing matrix S and 2) that this is the largest lower bound since $\tilde{K}(\varepsilon)$ achieves this in the limit. Thus, the infimum is $\|CS_*K\|_\infty$, with $K_{i,i} = \sum_{k=1}^n |S_*^{-1}B|_{i,k}$.□

The lemma suggests a general numerical method for finding the best upper bound in Theorem 1; identify a matrix S_* together with its inverse that diagonalize A and then apply our lemma with $B = C = I$. This is possible since effective tools exist for finding eigenvectors and matrix inverses. We also want to point out that, to the best of our knowledge, this result is not found in the linear algebra literature. An analogous result holds in the case of $\|\cdot\|_1$.

4. Case Study: PI-Control of Vehicle Formation

Here, we consider a network of double-integrator systems, where each agent is perturbed by a disturbance $w(t)$, i.e., $\ddot{x} = u(x) + w$. When restricted to only using relative feedback, a natural choice of controller is a second-order consensus system; both conventional and serial consensus can work. However, independent of the choice of second-order consensus system, the closed loop cannot reject a load disturbance $w(t) = w_0$ without stationary error. This can be seen by the fixed point conditions $L_{\text{pos}}x = w_0 \neq 0$ and $p_2p_1L^2x = w_0 \neq 0 \implies Lx \neq 0$ respectively. For measurable disturbances and with perfectly known dynamics, it is possible to design a feedforward term to reject any disturbance. Alternatively, an integral feedback term can be incorporated to handle unmeasured disturbances. In [Tegling et al., 2023], it was shown that distributed third-order conventional consensus may lose stability as more agents are added to the network. On the other hand, serial consensus is stable for any order [Hansson and Tegling, 2023]. We will, therefore, focus on the third-order serial consensus to handle potential load disturbances in the system. Following the serial-consensus design, the following feedback law can be utilized

$$\begin{aligned} u(x) = &-\underbrace{(p_1+p_2+p_3)}_{a_{\text{v}}} L(\dot{x} - v_{\text{ref}}\mathbf{1}) \\ &- \underbrace{(p_1p_2+p_1p_3+p_2p_3)}_{a_{\text{p}}} L^2(x-d) - \underbrace{(p_1p_2p_3)}_{a_{\text{I}}} Lz \\ \dot{z} = &\, L^2(x-d). \end{aligned} \tag{7}$$

The closed-loop system can then be analyzed in terms of the states $\dot{x} - v_{\text{ref}}\mathbf{1}$, $L(x-d)$, and z, which has the form

$$\begin{bmatrix} \dot{z} \\ L\frac{\mathrm{d}}{\mathrm{d}t}(x-d) \\ \frac{\mathrm{d}}{\mathrm{d}t}(\dot{x} - v_{\text{ref}}\mathbf{1}) \end{bmatrix} = \begin{bmatrix} 0 & L & 0 \\ 0 & 0 & L \\ -a_{\text{I}}L & -a_{\text{p}}L & -a_{\text{v}}L \end{bmatrix} \begin{bmatrix} z \\ L(x-d) \\ \dot{x} - v_{\text{ref}}\mathbf{1} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ w \end{bmatrix}. \tag{8}$$

This is of the same form as described in Theorem 1, so it will also exhibit scalable performance in terms of the transient deviation from any initial condition.

4.1 Disturbance rejection

Since only relative feedback is used, it is impossible to guarantee that any load disturbance can be rejected; if the graph underlying L contains a directed spanning tree, then any disturbance lying in $\text{Ker}(L) = \{v \in \mathbb{R}^N \mid v = a\mathbf{1},\ a \in \mathbb{R}\}$ will be undetectable by the controller. By the internal model principle [Francis and Wonham, 1976], it is therefore impossible for the controller to reject such a disturbance. For instance, the response to the constant disturbance $w = \mathbf{1}$ with all states starting at the origin will have the unbounded response $1^\top \dot{x}(t) = Nt$. But, if the disturbance lies in the image of L, or, if we only consider the relative states, the disturbance will be asymptotically rejected as per the following theorem.

THEOREM 2
The third-order serial consensus system (8) with distinct $p_k > 0$, $w(t) = Lw_0$, where w_0 is constant, and the graph underlying L contains a directed spanning tree, achieves $L(x-d) \to 0$ and $L\dot{x} \to 0$ as $t \to \infty$ and satisfy

$$\sup_{t\geq 0} \|\xi(t)\|_\infty \leq \alpha_\xi \|\xi(0)\|_\infty + \alpha_w \|w_0\|_\infty,$$

with the constants $\alpha_\xi = \|S_1\|_\infty \|S_1^{-1}\|_\infty$ and $\alpha_w = \frac{2}{p_1 p_2 p_3}\|S_2\|_\infty \|S_2^{-1}[1,0,0]^\top\|_\infty$ where S_1, S_2 are any matrices that diagonalize A.

The proof is found in the appendix. This theorem shows that the PI-controlled vehicle formation achieves scalable performance for the disturbances $z(0)$ and w_0, since $\|[e_{\rm p}^\top(t), e_{\rm v}^\top(t)]^\top\|_\infty \leq \alpha_\xi \|[e_{\rm p}^\top(0), e_{\rm v}^\top(0)]^\top\|_\infty + \alpha_\xi \|z(0)\|_\infty + \alpha_w \|w_0\|_\infty$. Furthermore, load disturbances will not lead to any stationary error. We remark that α_w and α_ξ can be optimized by applying Lemma 1.

5. Numerical Examples

First, we give a simple example of how to apply Lemma 1 and proceed with two vehicle formation examples.

5.1 Minimizing norm

EXAMPLE 1
Here, we want to find the smallest upper bound to the transient of the second-order serial consensus system $(sI + p_2 L)(sI + p_1 L)X = U$. A companion matrix A associated with $(s+p_1)(s+p_2)$ is

$$A = \begin{bmatrix} 0 & 1 \\ -p_1 p_2 & -(p_1+p_2) \end{bmatrix}.$$

Since we know that its eigenvalues are $-p_1, -p_2$ it is easy to find the eigenvectors: $s_1 = [1, -p_1]^\top$ and $s_2 = [1, -p_2]^\top$. We can define a diagonalizing matrix S_* and the corresponding inverse as

$$S_* = \begin{bmatrix} 1 & 1 \\ -p_1 & -p_2 \end{bmatrix}, \; S_*^{-1} = \frac{1}{p_1 - p_2}\begin{bmatrix} -p_2 & -1 \\ p_1 & 1 \end{bmatrix}.$$

Choosing $K = \frac{1}{|p_1 - p_2|}\mathrm{diag}(p_2+1, p_1+1)$ ensures that all rows of $K^{-1}S_*^{-1}$ have absolute sums 1. Finally, we calculate

$$\begin{aligned} \|S_* K\|_\infty &= \frac{\max(p_2 + p_1 + 2, p_1 p_2 + p_1 + p_2 + p_2 p_1)}{|p_1 - p_2|} \\ &= \frac{p_2 + p_1 + 2\max(1, p_1 p_2)}{|p_1 - p_2|}. \end{aligned}$$

This bound is also identical to the upper bound found in [Hansson and Tegling, 2024]. For a specific A, it is, of course, possible to perform all the previous steps with numerical methods instead.

5.2 PI control of vehicle formation

Example 2

For this example, we will consider a vehicular formation consisting of N identical double integrator systems $\ddot{x}_i = u_i(x)$. A directed path topology where each vehicle only observes its predecessors is used for the control. We refer back to Figure 1 for an illustration of the required communication structure. As illustrated, a signaling layer between neighboring vehicles is needed for the agents to use $L^2(x-d) = Le_{\mathrm{p}}$ and Lz. The associated graph Laplacian is

$$L_{\mathrm{ahead-path}} = \begin{bmatrix} 0 & 0 & & \\ -1 & 1 & & \\ & \ddots & \ddots & \\ & & -1 & 1 \end{bmatrix}.$$

The serial consensus-based control law (7) with $p_1 = 3$, $p_2 = 1$, and $p_3 = 1/3$ and $L = L_{\mathrm{ahead-path}}$ is used. In Figure 2, a simulation of when the lead vehicle starts at a desired reference velocity v_{ref} while all other vehicles start at rest is shown for 10 and 40 vehicles, respectively. A small transient error is observed before the vehicles return to the desired formation, as predicted by Theorem 2. In Figure 2c and Figure 2d, we have included a theoretical bound for the maximum transient relative error. While the bound $\|\xi(t)\|_\infty \leq \|S\|_\infty \|S^{-1}\|_\infty \|\xi(0)\|_\infty = 7 \cdot 10 = 70$ shows that no state will experience an unbounded transient, it is possible to improve this result. By replacing S with $[0,1,0]S$ and S^{-1} with $S^{-1}[0,0,1]^\top$, we isolate the transient effect from initial velocities on relative deviations. Additionally, by considering the state $\dot{x} - \mathbf{1}v_{\mathrm{ref}}/2$ instead of $\dot{x} - \mathbf{1}v_{\mathrm{ref}}$ the initial error will be half as big in the new frame. Applying the optimal bound from Lemma 1 yields a much tighter bound: $\|\xi(t)\|_\infty \leq \|[0,1,0]S\|_\infty \|S^{-1}[0,0,1]^\top\|_\infty \|\xi(0)\|_\infty = 1.5 \cdot 5 = 7.5$.

5.3 Vehicle formation with load disturbance

Example 3

In this example, we again consider a vehicle formation consisting of N identical double integrator systems with one virtual leader driving at a constant velocity. The vehicles drive along a road with constant inclination and thus experience a decelerating force (a load disturbance) due to gravity. The individual dynamics are $\ddot{x}_i = u_i(x,t) - g\frac{\theta}{\sqrt{1+\theta^2}}$, where $g = 9.8\,m/s^2$, is the gravitational acceleration constant and $\theta = 0.1$ is the inclination ratio. A second-order serial consensus simulation, with $L = L_{\mathrm{ahead-path}}$, $p_1 = 1/3$ and $p_2 = 3$ is shown in Figure 3a and Figure 3c. Here, a steady-state error exists due to the load disturbance. The third-order case, using the controller (7) with $L = L_{\mathrm{ahead-path}}$, $p_1 = 1/3$, $p_2 = 3$, and $p_3 = 1$ is shown

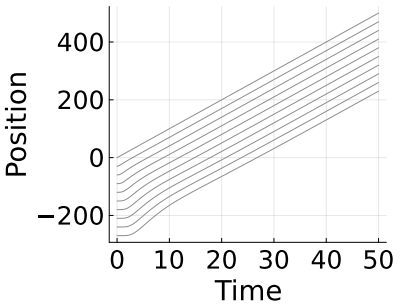

(a) Positions with $N = 10$ vehicles.

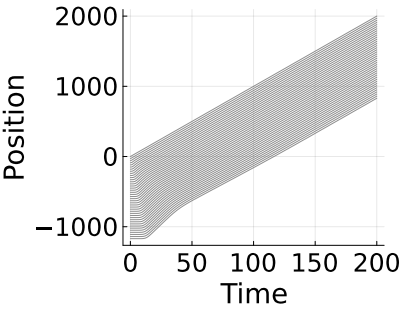

(b) Positions with $N = 40$ vehicles.

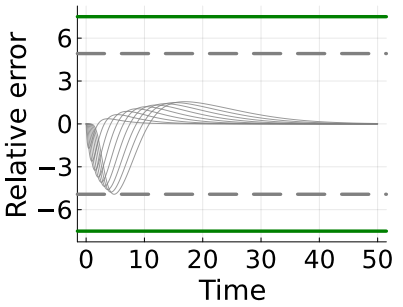

(c) Relative deviations from desired formation with $N = 10$. Dashed lines represent maximum transient relative error $\sup_{t\geq 0}\|e_{\mathrm{p}}\|_\infty = 4.92$ and solid green lines mark a theoretical bound $\sup_{t\geq 0}\|e_{\mathrm{p}}\|_\infty \leq 7.50$.

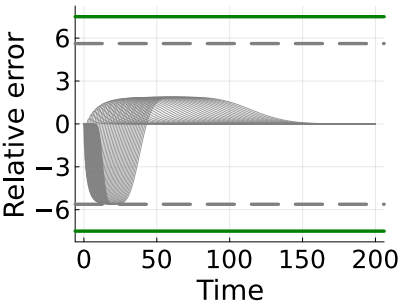

(d) Relative deviations from desired formation with $N = 40$. Dashed lines represent maximum transient relative error $\sup_{t\geq 0}\|e_{\mathrm{p}}\|_\infty = 5.62$ and solid green lines mark a theoretical bound $\sup_{t\geq 0}\|e_{\mathrm{p}}\|_\infty \leq 7.50$.

Figure 2. Simulation of lead vehicle driving at velocity $v_{\mathrm{ref}} = 10$ m/s and all other vehicles starting from a stand-still. Since the system has scalable performance, the transient error due to any initial condition will be bounded independently of the number of agents.

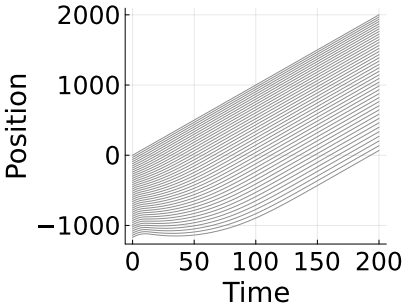

(a) Position for all vehicles, using proportional controller.

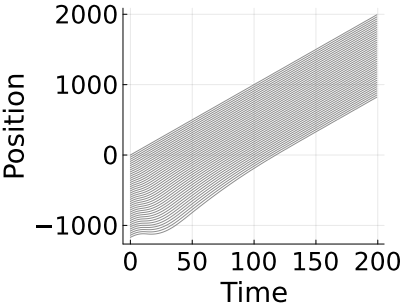

(b) Position for all vehicles, using PI controller.

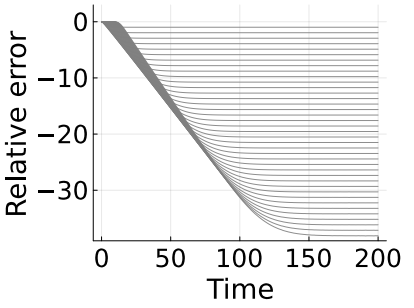

(c) Deviation from desired relative position for all vehicles, using proportional controller.

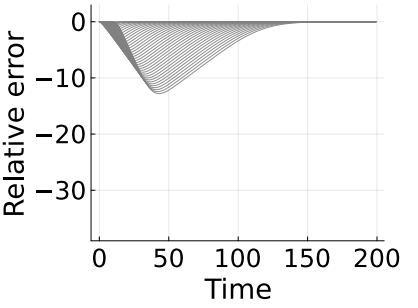

(d) Deviation from desired relative position for all vehicles, using PI controller.

Figure 3. Simulation of 40 agents driving uphill at an inclination ratio $\theta = 0.1$, and with a leader velocity 10 m/s. The agents that use a serial consensus-based PI controller experience a transient and then return to the desired spacing. The vehicles using a proportional controller experience a transient before settling with stationary error.

in Figure 3b and Figure 3d. In this case, a transient occurs before the vehicles eventually return to the desired formation. Unlike the initial value response, there is no uniform bound independent of the number of agents N for the transient response in this case. In fact, the bound of Theorem 2 requires that we identify a w_0 *that solves* $L_{\text{ahead}-\text{path}} w_0 = [0, \mathbf{1}^\top]^\top \frac{g\theta}{\sqrt{1+\theta^2}} \implies w_0 = \mathbf{1}a + [1, 2, \ldots, N]^\top \frac{g\theta}{\sqrt{1+\theta^2}}$. *Since this disturbance scales unboundedly with N, we cannot give a uniform upper bound for this disturbance.*

6. Discussion and Future Directions

This work has expanded the treatment of the newly proposed serial consensus protocol. The derived performance bound also suggests a simple numerical procedure for quantifying its performance.

PI control of vehicle formations serves as an exciting example of how third-order serial consensus can be used to reject load disturbances. Noteworthy, the additional integrator does not fundamentally impact the formation's performance. Our results also extend to fourth and higher-order consensus protocols, with limited practical examples. One compelling example is the coordination between platoons.

Although it seems that scalable input-to-state stability is not achievable through serial consensus without absolute feedback, a close notion of scalable performance is. We can bound the maximum transient response due to initial conditions and specific load disturbances, while the effects of more general disturbances remain an interesting direction for further study.

For future work, it would be interesting to evaluate how measurements such as GPS can be incorporated to achieve stronger notions of scalable performance, such as scalable input-to-state stability. The impact of implementation limitations, such as time delays, should also be addressed.

Acknowledgements

The authors are with the Department of Automatic Control and the ELLIIT Strategic Research Area at Lund University, Lund, Sweden. This work was partially funded by Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation and the Swedish Research Council through Grant 2019-00691.

References

Bamieh, B., M. R. Jovanovic, P. Mitra, and S. Patterson (2012). "Coherence in large-scale networks: dimension-dependent limitations of local feedback". *IEEE Trans. Autom. Control* **57**:9, pp. 2235–2249.

Besselink, B. and S. Knorn (2018). "Scalable input-to-state stability for performance analysis of large-scale networks". *IEEE Control Syst. Lett.* **2**:3, pp. 507–512.

Fax, J. A. and R. M. Murray (2004). "Information flow and cooperative control of vehicle formations". *IEEE Trans. Autom. Control* **49**:9, pp. 1465–1476.

Feng, S., Y. Zhang, S. E. Li, Z. Cao, H. X. Liu, and L. Li (2019). "String stability for vehicular platoon control: definitions and analysis methods". *Annu. Rev. Control* **47**, pp. 81–97.

Francis, B. A. and W. M. Wonham (1976). "The internal model principle of control theory". *Automatica* **12**:5, pp. 457–465.

Hansson, J. and E. Tegling (2023). "A closed-loop design for scalable high-order consensus". In: *IEEE Conf. Decis. Control*, pp. 7388–7394.

Hansson, J. and E. Tegling (2024). "Closed-loop design for scalable performance of vehicular formations". *arXiv:2402.15208*.

Jadbabaie, A., J. Lin, and A. S. Morse (2003). "Coordination of groups of mobile autonomous agents using nearest neighbor rules". *IEEE Trans. Autom. Control* **48**:6, pp. 988–1001.

Macellari, L., Y. Karayiannidis, and D. V. Dimarogonas (2017). "Multi-agent second order average consensus with prescribed transient behavior". *IEEE Trans. Autom. Control* **62**:10, pp. 5282–5288.

Olfati-Saber, R. and R. M. Murray (2004). "Consensus problems in networks of agents with switching topology and time-delays". *IEEE Trans. Autom. Control* **49**:9, pp. 1520–1533.

Pates, R., C. Lidström, and A. Rantzer (2017). "Control using local distance measurements cannot prevent incoherence in platoons". In: *IEEE Conf. Decis. Control*, pp. 3461–3466.

Patterson, S. and B. Bamieh (2014). "Consensus and coherence in fractal networks". *IEEE Trans. Control Netw. Syst.* **1**:4, pp. 338–348.

Ploeg, J., N. van de Wouw, and H. Nijmeijer (2014). "Lp string stability of cascaded systems: application to vehicle platooning". *IEEE Trans. Control Syst. Technol.* **22**:2, pp. 786–793.

Ren, W., K. L. Moore, and Y. Chen (2007). "High-order and model reference consensus algorithms in cooperative control of multi-vehicle systems". *J. Dyn. Syst. Meas. Control* **129**:5, pp. 678–688.

Ren, W. and E. Atkins (2005). "Second-order consensus protocols in multiple vehicle systems with local interactions". *Proc. AIAA GNC Conf.* **5**.

Siami, M. and N. Motee (2014). "Systemic measures for performance and robustness of large-scale interconnected dynamical networks". In: *IEEE Conf. Decis. Control*, pp. 5119–5124.

Stüdli, S., M. Seron, and R. Middleton (2017a). "From vehicular platoons to general networked systems: string stability and related concepts". *Annu. Rev. Control* **44**, pp. 157–172. ISSN: 1367-5788.

Stüdli, S., M. M. Seron, and R. H. Middleton (2017b). "Vehicular platoons in cyclic interconnections with constant inter-vehicle spacing". In: *IFAC-PapersOnLine*. Vol. 50. 1. Elsevier, pp. 2511–2516.

Tegling, E., B. Bamieh, and H. Sandberg (2023). "Scale fragilities in localized consensus dynamics". *Automatica* **153**, p. 111046.

Paper IV

Compositional Design for Time-Varying and Nonlinear Coordination

Jonas Hansson **Emma Tegling**

Abstract

This work addresses the design of multi-agent coordination through high-order consensus protocols. While first-order consensus strategies are well-studied—with known robustness to uncertainties such as time delays, time-varying weights, and nonlinearities like saturations—the theoretical guarantees for high-order consensus are comparatively limited. We propose a compositional control framework that generates high-order consensus protocols by serially connecting stable first-order consensus operators. Under mild assumptions, we establish that the resulting high-order system inherits stability properties from its components. The proposed design is versatile and supports a wide range of real-world constraints. This is demonstrated through applications inspired by vehicular formation control, including protocols with time-varying weights, bounded time-varying delays, and saturated inputs. We derive theoretical guarantees for these settings using the proposed compositional approach and demonstrate the advantages gained compared to conventional protocols in simulations.

Prepared for journal submission. Preprint available: J. Hansson and E. Tegling (2025). "Compositional design for time-varying and nonlinear coordination". arXiv: 2504.07226.

1. Introduction

Multi-agent coordination is one of the central problems in networked and distributed control. Consensus-seeking in opinions was modeled early on in [DeGroot, 1974] in a discrete-time setting, while [Levine and Athans, 1966; Melzer and Kuo, 1971; Chu, 1974a; Chu, 1974b] dealt with the coordination of vehicle strings. The problem was revisited in the early 2000's where a significant research thrust was initiated after some seminal works [Fax and Murray, 2004; Olfati-Saber and Murray, 2004; Olfati-Saber, 2006; Spanos et al., 2005; Moreau, 2005; Jadbabaie et al., 2003; Ren and Beard, 2005]. These established many of the fundamental properties of second- and first-order consensus protocols. Based on these works, we know that the first-order consensus protocol $\dot{x} = -L(t)x$, where $L(\cdot)$ is a time-varying graph Laplacian that encodes relative feedback, is robust to delays and time-varying topology. Furthermore, that consensus has a wide range of applications, ranging from swarming robots, vehicle platoons, frequency synchronization, and describing natural flocking behaviors [Olfati-Saber et al., 2007].

In this work, we will consider coordination among higher-order agents. If the first-order consensus protocol is $\dot{x} = -Lx$, the second-order is $\ddot{x} = -L_{\text{vel}}\dot{x} - L_{\text{pos}}x$, then there is a natural generalization to a general n^{th}-order consensus protocol, which is $x^{(n)} = -L_{(\text{n}-1)}x^{(n-1)} - \cdots - L_{(0)}x$. Here, the goal is to coordinate in position, velocity, and the high-order derivatives. This high-order consensus problem was first considered in [Ren et al., 2006; Ren et al., 2007], where some sufficient conditions for system stability were also established.

Our motivation for revisiting this problem comes from studies of the dynamic behaviors of, in particular, large-scale coordinating multi-agent networks. Vehicular platoons suffering from error propagation and string stability challenges [Yadlapalli et al., 2006; Stüdli et al., 2017a], fundamental limitations in terms of formation coherence [Bamieh et al., 2008; Bamieh et al., 2012], and lately scale fragilities [Tegling et al., 2023; Tegling et al., 2019] in second- and higher-order consensus. Here, poor dynamical behaviors—or even instability—can emerge when networks grow large in a manner that is hard to foresee from the original, distributed, control design. This calls for methods to construct coordination protocols that allow for a modular and scalable network design. An early and influential such approach based on passivity was [Arcak, 2007]; here, we take an alternative route.

Apart from challenges related to large-scale dynamic behaviors, even the problem of stabilization is non-trivial in higher-order coordination, and more difficult than in first- and second-order protocols. For instance, the first-order protocol permits time-varying topologies with certain time-delayed measurements [Lu and Liu, 2017], saturations [Li et al., 2011], and directed topologies provided the network is sufficiently connected over time. For the second-order linear and time-invariant consensus protocol, stability can be guaranteed, provided that $L_{\text{vel}} = r_{\text{vel}}L$ and $L_{\text{pos}} = r_{\text{pos}}L$, where L is symmetric and contains a spanning tree. When the symmetry condition is broken, as in the case of a directed cycle graph, then the second-order linear consensus

protocol can become unstable [Stüdli et al., 2017b]. There are also many works that have derived sufficient conditions for this protocol when subject to time-delays [Gao et al., 2023], time-varying structures [Li et al., 2022], and various nonlinearities [Lyu et al., 2016]. However, these results most often depend on global or absolute knowledge of the positions and velocities. Higher-order protocols have been shown to lack scalable stability in sparse networks [Tegling et al., 2023], meaning that a loss of closed-loop stability is inevitable without a (re-) tuning based on global knowledge. Other works on high-order coordination include [Trindade et al., 2024] that studied LQR, [Li et al., 2024; Tian and Zhang, 2012] time-varying topology and delays, and [Rezaee and Abdollahi, 2015; He and Cao, 2011; Liu and Jia, 2010], where consensus is achieved, but with the help of absolute feedback.

In this work, we propose a novel consensus protocol for achieving coordination in a network of n^{th}-order integrators. Our proposed control design is based on the idea of first designing the closed-loop system and then identifying the corresponding control law. The class of desired closed-loop systems can be written as the composition of n simple first-order consensus systems, that is,

$$\left(\frac{\mathrm{d}}{\mathrm{d}t}+\mathcal{L}_n\right)\circ\cdots\circ\left(\frac{\mathrm{d}}{\mathrm{d}t}+\mathcal{L}_1\right)(x)=0,$$

where each $\mathcal{L}$ describes a, potentially nonlinear and time-varying operator that generalizes the graph Laplacian in the LTI case. Due to its compositional structure, we will call this the *compositional consensus* system. In the second-order case, this can be expanded to

$$\ddot{x}=u(x,t)=-\mathcal{L}_2(\dot{x}+\mathcal{L}_1(x,t),t)-\frac{\mathrm{d}}{\mathrm{d}t}\mathcal{L}_1(x,t).$$

Under relatively mild conditions on the operators $\mathcal{L}_k$, essentially that their corresponding first-order protocols achieve consensus, we can show that this control design guarantees that the solution x, $\dot{x}$ and all its first $n-1$ derivatives will coordinate and reach an n^{th}-order consensus. We note that this is independent of the underlying graph structure.

To illustrate the strength of the compositional consensus, we also demonstrate how to apply this controller when the composing components $\mathcal{L}_k$ correspond to saturated, time-varying, and time-delayed consensus protocols, building on and partially extending results existing in literature. In particular, we prove a general result on the stability of consensus under saturated control inputs. We formally and through case studies show that compositional consensus has superior stability and performance than a more naïve higher-order protocol. The implementation of the protocol remains localized, but requires some additional signaling in an $n-$hop neighborhood, or message-passing between nearest neighbors. Due to the strong robustness towards time-delays and time-varying connectivity, such signaling need not be ideally implemented.

Paper Outline The remainder of the paper is organized as follows. We next introduce some notation and preliminaries, followed by an introduction of compositional consensus in Section 2. Section 3 presents our main result along with key lemmas used in the proof. Section 4 studies some selected first-order consensus protocols that can be used in the compositional design. In Section 5 we illustrate our result through numerical simulations, and Section 6 concludes the paper.

1.1 Mathematical preliminaries

Graph theory We represent a directed graph as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{1, \dots, N\}$ is the set of nodes, and $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ is the set of edges. The graph is associated with a weighted adjacency matrix $W \in \mathbb{R}^{N \times N}$, where $W_{i,j} > 0$ if and only if $(j,i) \in \mathcal{E}$, i.e., agent j influences agent i. The corresponding graph Laplacian is defined as

$$L = D - W,$$

where D is the diagonal degree matrix with $D_{i,i} = \sum_{j=1}^{N} W_{i,j}$. The graph Laplacian L has zero row sum and encodes the relative feedback structure of the network.

A graph is said to contain a *directed spanning tree* if there exists a node $k \in \mathcal{V}$ such that all other nodes $j \in \mathcal{V} \setminus \{k\}$ are reachable via a directed path from k. If this condition holds, the Laplacian L has a simple zero eigenvalue, and all other eigenvalues have strictly positive real parts.

We also make use of δ*-graphs* as defined in [Moreau, 2004]. Given a threshold $\delta > 0$, the δ-graph associated with W is a subgraph where an edge (j,i) is retained if and only if $W_{i,j} \geq \delta$.

Norms and other notation We denote vector and matrix norms using $\|\cdot\|$. For vectors $x \in \mathbb{C}^N$, we use the ∞-norm

$$\|x\|_\infty = \max_i |x_i|,$$

and for matrices $C \in \mathbb{C}^{N \times N}$, the induced matrix norm

$$\|C\|_\infty = \max_i \sum_j |C_{ij}|.$$

Seminorms are denoted $|||\cdot|||$, following the notation in [Bullo, 2024]. These are functions satisfying the triangle inequality $|||x_1 + x_2||| \leq |||x_1||| + |||x_2|||$ and absolute homogeneity $|||ax||| \leq |a|\,|||x|||$. When the context is clear, we drop the explicit time-dependence $x(t)$ in the notation. We write $\frac{\mathrm{d}^j}{\mathrm{d}t^j} x = x^{(j)}$ for the j^{th} time derivative, and use $\partial_t F(x,t)$ for partial derivatives.

Function composition is written $(f \circ g)(x) = f(g(x))$. When composing time-varying functions, we use the convention

$$(\mathcal{L}_2 \circ \mathcal{L}_1)(x) := \mathcal{L}_2(\mathcal{L}_1(x,t),t).$$

A continuous function $\gamma(\cdot)$ is said to be of class $\mathcal{K}$ if it is non-negative and strictly increasing. A continuous function $\beta(\cdot,\cdot)$ belongs to class $\mathcal{KL}$ if it for any fixed s, $\beta(\cdot,s)$ is of class $\mathcal{K}$, and for any fixed r, $\beta(r,\cdot)$ is decreasing with respect to s and satisfy $\lim_{s\to\infty}\beta(r,s)=0$. This follows the standard notation of [Khalil, 2002].

1.2 Consensus

Due to the presumed lack of absolute feedback (see Assumption 1), the relevant notion of stability in this work is instead one of consensus among the agents. It is defined as follows.

Definition 1—Consensus
Let $x(t)\in\mathbb{R}^N$ be the state of a multi-agent system governed by $\dot{x}=f(x,t)$. The system is said to achieve consensus *if*

$$\lim_{t\to\infty}|x_i(t)-x_j(t)|=0, \quad \textit{for all } i\neq j.$$

It is well known that the simple, linear, continuous-time consensus protocol

$$\dot{x}=-Lx,$$

achieves consensus, where L is a graph Laplacian, provided the underlying graph contains a directed spanning tree [Lin et al., 2005; Ren et al., 2005]. In high-order coordination problems, synchronizing the positions and higher-order derivatives, such as velocities and accelerations, is often desirable. This motivates the following generalization (see also [Ren et al., 2006]):

Definition 2—n^{TH}-Order consensus
Let $x(t)\in\mathbb{R}^N$ evolve according to

$$\frac{\mathrm{d}^n x}{\mathrm{d}t^n}=f(x,t).$$

The solution x is said to achieve n^{th}-order consensus if

$$\lim_{t\to\infty}|x_i^{(k)}(t)-x_j^{(k)}(t)|=0,\ \forall i\neq j \textit{ and } k=0,\ldots,n-1.$$

This definition captures the idea that all agents eventually align in their positions and higher-order dynamics, like velocities and accelerations.

2. Problem Setup

In this work we consider a network consisting of N identical agents with n^{th}-order integrator dynamics, that is,

$$x^{(n)}(t)=u(x,t), \tag{1}$$

$x(0) = x_0,\ \dot{x}(0) = \dot{x}_0, \ldots,\ x^{(n)}(0) = x_0^{(n)}$. Our proposed control design can be compared with a Youla-Kucera parametrization, where we first design the closed-loop system to be

$$\left(\frac{\mathrm{d}}{\mathrm{d}t} + \mathcal{L}_n\right) \circ \cdots \circ \left(\frac{\mathrm{d}}{\mathrm{d}t} + \mathcal{L}_2\right) \circ \left(\frac{\mathrm{d}}{\mathrm{d}t} + \mathcal{L}_1\right)(x) = 0, \tag{2}$$

where each operator $\mathcal{L}_k(\cdot, t) : \mathbb{R}^N \to \mathbb{R}^N$ is allowed to be time-dependent and potentially nonlinear. Our controller is then chosen to be the one that achieves this closed-loop system, that is,

$$u(x,t) = x^{(n)}(t) - \left(\frac{\mathrm{d}}{\mathrm{d}t} + \mathcal{L}_n\right) \circ \cdots \circ \left(\frac{\mathrm{d}}{\mathrm{d}t} + \mathcal{L}_1\right)(x). \tag{3}$$

The closed-loop design matches the behavior of n dynamical systems in a series interconnection. The system can be analyzed in the following simple state-space form

$$\begin{bmatrix} \dot{\xi}_1 \\ \vdots \\ \dot{\xi}_{n-1} \\ \dot{\xi}_n \end{bmatrix} = \begin{bmatrix} -\mathcal{L}_1(\xi_1, t) + \xi_2 \\ \vdots \\ -\mathcal{L}_{n-1}(\xi_{n-1}, t) + \xi_n \\ -\mathcal{L}_n(\xi_n, t) \end{bmatrix}. \tag{4}$$

where $\xi_1 = x$ and $\xi_{k+1} = \dot{\xi}_k + \mathcal{L}_k(\xi_k, t)$. We want to emphasize that the state-space formulation (4) is the key to the scalability of the compositional consensus formulation. The series interconnection of dynamical systems has some favorable properties. For instance, any series interconnection of strongly contracting systems will itself be strongly infinitesimally contracting [Bullo, 2024, Theorem 3.23]. We also want to highlight the connection to the literature on distributed optimization, using gradient tracking [Carnevale et al., 2023; Dhullipalla and Chen, 2024], and also used for dynamic average consensus [Kia et al., 2019; Aldana-López et al., 2022]. One key difference between the mentioned work and ours is that we will mainly focus on the scenario where $\mathcal{L}_k$ are restricted to using only *relative feedback*. A limitation to relative feedback poses severe challenges in coordination control design, see e.g. [Bamieh et al., 2012], and is motivated by a fundamental difficulty in many applications to capture absolute position, phase, etc., whereas the corresponding relative measurement is readily available. The following Assumption captures this limitation.

Assumption 1—Relative feedback
The feedback operators satisfy $\mathcal{L}_k(z(t) + \mathbf{1}a(t), t) = \mathcal{L}_k(z(t), t)$ *for any* z, $a(\cdot)$, *and* t.

Under suitable and relatively simple conditions, it is possible to show that the solution $x(t)$ of the closed-loop system (2) will converge to an n^{th}-order consensus. Furthermore, that sparsity of the individual operators $\mathcal{L}_k$ implies sparsity of the controller $u(x,t)$, as defined by (3).

Example 1
In the linear, time-invariant case, the composition (2) may capture the serial consensus *protocol. Here, the closed loop matches regular consensus protocols connected in a series. The n^{th}-order serial consensus system can, in the Laplace domain, be represented as*

$$(sI+L_n)\cdots(sI+L_2)(sI+L_1)X(s)=U_{\text{ref}}(s).$$

One property that makes this system interesting is the simple condition for stability. That is, if each of the graphs underlying the L_k's contain a directed spanning tree, then this high-order consensus protocol will achieve an n^{th}-order consensus, assuming a decaying input signal $\|u_{\text{ref}}(t)\|\to 0$ [Hansson and Tegling, 2023]. The serial consensus protocol can also be used to construct linear time-invariant consensus protocols for vehicular formations with a strong notion of scalable performance [Hansson and Tegling, 2025; Hansson and Tegling, 2024]. It, therefore, avoids issues with scale fragility [Tegling et al., 2023] and string instability [Stüdli et al., 2017a] affecting conventional consensus protocols.

To implement serial consensus, additional signaling may be needed in the multi-agent system. This can be seen in the second-order serial consensus where the control law is

$$u(x,t)=-(L_1+L_2)\dot{x}-L_2L_1x.$$

Here, the velocity feedback can be realized immediately through local measurements. For the second term, each agent can first aid in calculating $e=L_1x$, then message pass this measurement to their followers so that they can compute the relative differences L_2e. In general, it is possible to compute the local control law for the n^{th}-order serial consensus through the use of $n-1$ local message passes; the local consensus protocol is only dependent on relative measurements within an n-hop neighborhood (at most) of each agent.

2.1 Assumptions

To prove the main result, we will use the following assumptions, which apply to the operators $\mathcal{L}_k$ for $k\leq n-1$. First, we impose a standard technical assumption used to establish the existence and uniqueness of a solution.

Assumption 2
The $\mathcal{L}_k(z,t)$ is Lipschitz in z with a global Lipschitz constant independent of t and are, for any fixed z, piecewise continuous in t.

The next assumption is one of input-to-state stability (ISS) for the individual subsystems in the composition.

Assumption 3
If $\|w(t)\|\leq M_k$ for all $t\geq T_0$, then the system $\dot{z}=\mathcal{L}_k(z,t)+w(t)$ is ISS with respect to some seminorm $|||\cdot|||$, that is:

$$|||z(t)|||\leq\beta_k(|||z(T_0)|||,t)+\gamma_k(\sup_{t\geq T_0}\|w(t)\|)$$

where $\beta_k \in \mathcal{KL}$, $\gamma_k \in \mathcal{K}$, and the seminorm satisfies $|||z||| = 0 \iff z \in \text{span}(\mathbf{1})$.

It implies consensus of the individual subsystems and will be needed to prove consensus of the composed system. Finally,

Assumption 4
Let $\mathcal{L}_k \in C^{n-1-k}$ be chosen such that $\| \frac{\mathrm{d}^j}{\mathrm{d}t^j} \mathcal{L}_k(z,t) \| \leq \alpha_{k,j}(\max_{0 \leq i \leq j} \|z^{(j)}\|)$ for some functions $\alpha_{k,j} \in \mathcal{K}$, for all $j \leq n-k-1$, and all time $t \geq 0$.

This assumption asserts a smoothness of the composing operators $\mathcal{L}_k$. With this assumption, we can prove that coordination of ξ_k is equivalent to the coordination of $x, \dot{x}, \ldots, x^{n-1}$. With these assumptions established, we are ready to state our main theorems.

3. Main Results

Consider the following result, which establishes that the composition of consensus protocols according to (2) will also achieve consensus.

Theorem 1
Let each subsystem $\mathcal{L}_k$, implement relative feedback according to Assumption 1, and be chosen such that each unperturbed system

$$\dot{z}_k = \mathcal{L}_k(z_k, t)$$

admits a unique solution for every initial condition $z_k(0)$ that satisfies

$$\lim_{t \to \infty} \|z_k(t) - \mathbf{1}a_k(t)\| = 0,$$

for some function a_k. Assume additionally that each $\mathcal{L}_k$ satisfies Assumptions 2–4 for $k = 1, \ldots, n-1$. Then, the compositional consensus system (2) admits a unique solution x, and this solution achieves n^{th}-order consensus.

We now present the lemmas that form the basis of the proof of Theorem 1.

Lemma 1
If all $\mathcal{L}_k$, $k = 1, \ldots, n$ implement relative feedback according to Assumption 1, and satisfies the smoothness and boundedness Assumption 4 for all $k \leq n-1$, and all times $t \geq 0$. Then, the following two are equivalent:

i) The solution x of the compositional consensus (2) achieves n^{th}-order consensus for any initial condition;

ii) the states ξ_k, $k = 1, \ldots, n$, of (4) achieve first-order consensus for any initial condition.

Proof sketch. (Full proof given in Appendix A.) We prove the equivalence by induction in two directions. First, we show that the initial condition of x and its first $n-1$ derivatives uniquely determine the initial conditions of the states ξ_k. Then, using a similar argument, the reverse direction can be proven.

Since $x(t) = \xi_1(t)$, and Assumption 4 ensures sufficient smoothness, we may recursively differentiate this relation to recover all $\xi_k(t)$. The supporting Lemma 3, a consequence of the relative feedback Assumption 1, allows us to bound the terms $\frac{\mathrm{d}^j}{\mathrm{d}t^j}\mathcal{L}_k(\xi_k,t)$ in terms of deviations from consensus

$$\|\frac{\mathrm{d}^j}{\mathrm{d}t^j}\mathcal{L}_k(\xi_k,t)\| \leq \alpha_{k,j}(\max_{0\leq i\leq j}(\|\xi_k^{(i)} - \mathbf{1}b_{k+i}(t)\|).$$

We then apply induction in k to show that n^{th}-order consensus of x implies first-order consensus of all ξ_k and induction in j to show the converse. The full details are provided in Appendix A. □

Having established the equivalence between (2) and (4), we now show that the latter achieves consensus under relatively mild conditions.

Lemma 2
Let each subsystem $\mathcal{L}_k$ implement relative feedback according to Assumption 1, and assume that the unperturbed system

$$\dot{z}_k = \mathcal{L}_k(z_k,t)$$

admits a unique solution for any initial condition $z_k(0)$, and satisfies

$$\lim_{t\to\infty}\|z_k(t) - \mathbf{1}b_k(t)\| = 0$$

for some function b_k. Assume additionally that each $\mathcal{L}_k$ satisfies Assumptions 2 and 3 for $k = 1,\ldots,n-1$. Then, the states ξ_k in (4) admit a unique solution, and satisfy

$$\lim_{t\to\infty}\|\xi_k(t) - \mathbf{1}a_k(t)\| = 0$$

for some functions a_k.

Proof sketch. (Full proof given in Appendix B.) Existence and uniqueness follow from Carathéodory's existence and uniqueness theorem, which here follows from Assumption 2. The solution for ξ_n exists and achieves consensus by assumption. The remaining states ξ_k are shown to reach consensus through induction on k.

In particular, through Assumption 1 we establish that

$$\dot{\tilde{\xi}}_k = -\mathcal{L}_k(\tilde{\xi}_k,t) + w(t),$$

where $\tilde{\xi}_k(t) := \xi_k(t) - \mathbf{1}\int_0^t a_{k+1}(\tau)\,\mathrm{d}\tau$, and $w(t) := \xi_{k+1}(t) - \mathbf{1}a_{k+1}(t)$. The input $w(t)$ converges to zero due to the inductive hypothesis.

Then, by the local ISS property in Assumption 3, it follows that $\tilde{\xi}_k$ achieves first-order consensus, which in turn implies consensus of ξ_k. Repeating this argument recursively establishes the result for all ξ_k. □

With Lemmas 1 and 2, the main result in Theorem 1 is now readily established. Together, these lemmas provide sufficient conditions for the solution x to achieve nth-order consensus. Theorem 1 is thereby proven.

Remark 1
The existence and uniqueness of the solution x should be interpreted as a weak solution, i.e., a function that satisfies the differential equation (2) almost everywhere. If a smooth solution is desired—i.e., one that satisfies the equation pointwise—one may strengthen Assumption 4 by requiring $\mathcal{L}_k \in C^{n-k}$ instead of $\mathcal{L}_k \in C^{n-k-1}$.

3.1 Example: serial consensus

To better illustrate Theorem 1, consider again the simplest case where each operator is a linear time-invariant function, i.e., $\mathcal{L}_k(x,t) = L_k x$, which leads to the compositional consensus system also known as *serial consensus*:

$$\left(\frac{\mathrm{d}}{\mathrm{d}t} + L_n\right) \cdots \left(\frac{\mathrm{d}}{\mathrm{d}t} + L_1\right)(x) = 0. \tag{5}$$

We now show the following.

Proposition 1
If each graph Laplacian L_k in (5) contains a (possibly different) directed spanning tree, then x achieves n^{th}-order consensus for any initial condition.

The proof of this proposition serves as an example of how to apply our compositional consensus result.

*Example 1—**continued***
We verify the assumptions of Theorem 1 for the case $\mathcal{L}_k(x,t) = L_k x(t)$. First, since

$$\|L_k x - L_k y\| \leq \|L_k\| \|x - y\|,$$

each operator is Lipschitz and time-invariant, so Assumption 2 is satisfied. The invariance property of Assumption 1 follows from the definition of the Laplacian, since $L_k(\mathbf{1}a(t)) = 0$ for any scalar function $a(t)$.

The most involved step is verifying Assumption 3. We define the seminorm $|||z|||_k := \|L_k z\|$, which is valid since L_k has a simple zero eigenvalue with corresponding eigenvector $\mathbf{1}$ (by the spanning tree condition). Consider the perturbed consensus system

$$\dot{z}(t) = -L_k z(t) + w(t).$$

Its general solution is given by

$$z(t) = e^{-L_k(t-T_0)}z(T_0) + \int_{T_0}^{t} e^{-L_k(t-\tau)}w(\tau)\,\mathrm{d}\tau.$$

Premultiplying by L_k, taking norms, and using the bound $\|L_k e^{-L_k t}\| \leq M_k e^{-\alpha t}$, valid for some $M_k, \alpha > 0$, yields

$$\|L_k z(t)\| \leq M_k e^{-\alpha(t-T_0)}\|L_k^+\|\|L_k z(T_0)\| + \frac{M_k}{\alpha}\sup_{\tau \geq T_0}\|w(\tau)\|,$$

where L_k^+ denotes a pseudoinverse of L_k. From this inequality, one can identify $\beta_k(\cdot) \in \mathcal{KL}$ and $\gamma_k(\cdot) \in \mathcal{K}$, verifying the local ISS property. That the system $\dot{z} = -L_n z$ asymptotically reaches consensus is well known, but it is also a direct consequence of the above discussion.

Finally, Assumption 4 concerns the smoothness of $\mathcal{L}_k$. Since $\partial_t \mathcal{L}_k(x,t) = 0$, $\partial_x \mathcal{L}_k(x,t) = L_k$, and all higher-order derivatives

$$\frac{\partial^{i+j}}{\partial x^j \partial t^i}\mathcal{L}_k(x,t) = 0$$

vanish, it follows that $\mathcal{L}_k \in C^{n-1-k}$ as required. The time derivative of $\mathcal{L}_k(x,t)$ is

$$\frac{\mathrm{d}^j}{\mathrm{d}t^j}\mathcal{L}_k(x,t) = L_k x^{(j)},$$

so

$$\left\|\frac{\mathrm{d}^j}{\mathrm{d}t^j}\mathcal{L}_k(x,t)\right\| \leq \|L_k\|\|x^{(j)}\|,$$

which satisfies Assumption 4 with bounding functions $\alpha_{k,j}(r) = \|L_k\|r|$, which clearly are of class $\mathcal{K}$.

This result confirms the stability of the linear serial consensus system previously established by different methods in [Hansson and Tegling, 2023]. Although many assumptions need to be checked, most are straightforward. The more involved ones—Assumptions 3 and the stability of unperturbed first-order system—can be verified using classical first-order consensus theory, as we will demonstrate in the following applications.

3.2 Implementation of compositional consensus

The implementation of compositional consensus raises a few key questions: 1) Is the protocol implementable using only local and relative feedback? 2) Will the control signal be well-defined?

The answer to the first question is yes; provided that each $\mathcal{L}_k$ implements relative local feedback, the resulting feedback will also be relative and local. To illustrate this, consider the third-order case

$$x^{(3)} = -\mathcal{L}_3(\ddot{x}+\mathcal{L}_2(\dot{x}+\mathcal{L}_1(x,t),t)+\frac{\mathrm{d}}{\mathrm{d}t}\mathcal{L}_1(x,t),t) \\ -\frac{\mathrm{d}}{\mathrm{d}t}\mathcal{L}_2(\dot{x}+\mathcal{L}_1(x,t),t)-\frac{\mathrm{d}^2}{\mathrm{d}t^2}\mathcal{L}_1(x,t)$$

and suppose that the unweighted adjacency matrix W_k encodes the communication structure of $\mathcal{L}_k$, $k=1,2,3$. That is, $[W_k]_{i,j}=1 \iff [\mathcal{L}_k(z+\mathbf{e}_j,t)-\mathcal{L}_k(z,t)]_i \neq 0$ where $\mathbf{e}_j$ is the j^{th} unit vector. We may now work backward to deduce the adjacency matrix encoding the full feedback. The term $z_1=\mathcal{L}_1(x,t)$ depends on measurements coming from the graph associated with W_1 and so will all its higher derivatives. Let $z_2=\mathcal{L}_2(\dot{x}+z_1,t)$, which then depends on signals encoded by $W_2(W_1+I)$. Finally, since $z_3=\mathcal{L}_3(\ddot{x}+z_2+\dot{z}_1)$, these signals will be encoded by $W_3(I+W_2(I+W_1)+W_1)$. In general, we see that all measurements needed in the feedback are contained by the graph associated with $(W_k+I)(W_{k-1}+I)\cdots(W_1+I)$. In the special case where the W_k are identical, the product implies a k-hop neighborhood in the graph in question.

As to the second question, it is not in general guaranteed that the highest derivative $x^n(t)$ is well-defined for all t in the fairly general setting of Theorem 1. The issue can be effectively illustrated by considering the second-order case.

$$u(x,t) = -\mathcal{L}(\dot{x}+\mathcal{L}_1(x,t),t)-\frac{\mathrm{d}}{\mathrm{d}t}\mathcal{L}_1(x,t).$$

While the first term suffers no problem, the term $\frac{\mathrm{d}}{\mathrm{d}t}\mathcal{L}_1(x,t)$ may be problematic, since Theorem 1 only requires $\mathcal{L}_1 \in C^0$. The derivative term may, therefore, instead be interpreted in terms of the Dini derivative, that is,

$$D^+(\mathcal{L}_1(x(t),t) = \limsup_{\Delta t>0,\Delta t\to 0}\frac{\mathcal{L}_1(x(t+\Delta t),t+\Delta t)}{\Delta t}.$$

This will always be well-defined; see e.g. [Bullo, 2024, A.7], however, potentially unbounded. From a more practical view, one can consider a function that approximates the derivative almost everywhere. This is relevant when using nonlinear functions like the 1- and ∞-norms, saturations, and dead zones. In the case of saturations one may use $D^+\mathrm{sat}(x_i(t)) \overset{\text{a.e.}}{=} \dot{x}_i(t)\mathbf{I}_{(-1,1)}(x_i(t))$ where $\mathbf{I}_{(-1,1)}(\cdot)$ is an indicator function.

4. Applications of Compositional Consensus

In this section, we explore applications of Theorem 1 in representative nonlinear and time-varying networked systems. This amounts to verifying whether the protocols satisfy the Assumptions in Section 2.1.

4.1 Saturated consensus

A common type of nonlinearity in control systems is *saturation*, which arises due to actuator limitations or other physical constraints. This example shows that saturated signals can be handled within the compositional consensus framework.

Earlier works, such as [Li et al., 2011], have established asymptotic consensus stability for the unforced system

$$\dot{z} = -\mathrm{sat}(Lz).$$

This makes the corresponding operator admissible as the outermost function in the compositional consensus (2), i.e., $\mathcal{L}_n(z,t) = \mathrm{sat}(Lz)$. In the following, we extend the analysis to bounded-input scenarios, thereby enabling the use of saturation-based dynamics also for $\mathcal{L}_{n-1}(z,t)$; this follows since the remaining assumptions are simple to check. The following proposition proves the applicability to Assumption 3.

Proposition 2
Consider the consensus system

$$\dot{z} = -\mathrm{sat}(Lz) + d(t),$$

where L is a graph Laplacian that contains a directed spanning tree. Then there exists a constant $d_{\max} > 0$ such that, for all disturbances satisfying $\|d(t)\|_\infty < d_{\max}$, the disagreement satisfies

$$\|Lz(t)\|_\infty \leq \beta(\|Lz(0)\|_\infty, t) + \gamma(\sup_{t\geq 0} \|d(t)\|_\infty)$$

for some functions $\beta \in \mathcal{KL}$ and $\gamma \in \mathcal{K}$.

The proof is provided in Appendix C. This result is applied to vehicular coordination in Section 5.2.

4.2 Time-varying linear dynamics

Another class of systems that can be used within the compositional consensus framework is linear time-varying dynamics. The following proposition is a straightforward application of [Moreau, 2004, Theorem 1].

Proposition 3
Consider the first-order consensus system

$$\dot{z} = -L(t)z + d(t),$$

where $L(t)$ is a piecewise continuous Metzler matrix. Let $A(t) = \int_t^{t+T} L(\tau)\,\mathrm{d}\tau$. Suppose there exists $\delta > 0$, a fixed node k, and a time $T > 0$ such that, for every t, the

δ-digraph associated with $A(t)$ contains a node k that is reachable from all other nodes. Define the disagreement seminorm

$$|||z||| := \left\| \left(I - \frac{\mathbf{1}\mathbf{1}^\top}{N} \right) z \right\|.$$

Then, the solution satisfies

$$|||z(t)||| \leq \beta\left(|||z(0)|||, t\right) + \gamma(\sup_{t\geq 0}\|d(t)\|),$$

for some functions $\beta \in \mathcal{KL}$ and $\gamma \in \mathcal{K}$.

Proof. The cited theorem [Moreau, 2004, Theorem 1] establishes that the consensus equilibrium of the unperturbed system (i.e., $d(t) = 0$) is uniformly exponentially stable, meaning

$$|||z(t)||| \leq Me^{-\alpha t}|||z(0)|||$$

for some constants $M > 0$ and $\alpha > 0$. Let $\Phi(t, t_0)$ denote the state transition operator of the unforced system $\dot{z} = -L(t)z$. Then, by the variation of constants formula (see, e.g., [Hinrichsen and Pritchard, 2005, Chapter 2]), the solution to the forced system is

$$z(t) = \Phi(t, t_0)z(t_0) + \int_{t_0}^{t} \Phi(t, \tau)d(\tau)\,\mathrm{d}\tau.$$

Applying the projection operator $I - \frac{\mathbf{1}\mathbf{1}^\top}{N}$ to both sides, and using the fact that $\Phi(t, t_0)$ preserves the consensus subspace, we get

$$|||z(t)||| \leq Me^{-\alpha(t-t_0)}|||z(t_0)||| + \int_{t_0}^{t} Me^{-\alpha(t-\tau)}\|d(\tau)\|\,\mathrm{d}\tau.$$

Using the standard exponential convolution estimate, we obtain

$$|||z(t)||| \leq Me^{-\alpha(t-t_0)}|||z(t_0)||| + \frac{M}{\alpha}\sup_{t\geq 0}\|d(t)\|,$$

which is an ISS-type bound of the desired form. □

For $\mathcal{L}_k(z, t) = L_k(t)z$, it is straightforward to verify Assumptions 2–3. What remains is to establish the smoothness condition in Assumption 4. Note that

$$\begin{aligned}
\partial_t^j(L_k(t)z) &= L_k^{(j)}(t)z, \\
\partial_t^j\partial_z(L_k(t)z) &= L_k^{(j)}(t), \\
\partial_t^j\partial_z^2(L_k(t)z) &= 0.
\end{aligned}$$

This shows that $\mathcal{L}_k \in C^n$ if and only if $L_k(t)$ is n-times continuously differentiable.

Furthermore, applying the product rule yields

$$\frac{\mathrm{d}^j}{\mathrm{d}t^j}(L_k(t)z(t)) = \sum_{i=0}^{j} \binom{j}{i} L_k^{(i)}(t) z^{(j-i)},$$

which can be uniformly bounded in terms of $\max_{0\le i\le j}\left\{\|z^{(i)}\|\right\}$ provided that $\|L_k^{(i)}(t)\| \le M$ for all $0 \le i \le j$ and some constant $M > 0$.

Remark 2
Similar to the argument above, one may also apply [Khalil, 2002, Lemma 4.6] to the system $\dot{x} = -\mathcal{L}(z,t) + d(t)$ to establish that uniform exponential stability implies ISS, provided that $\mathcal{L}(z,t)$ is continuous.

4.3 Time-delayed consensus

As a final case, we consider consensus protocols with time delays, modeled by functional differential equations. These systems have been thoroughly examined in, e.g., [Moreau, 2004; Lu and Liu, 2017], which establish sufficient conditions for asymptotic consensus in the presence of bounded communication delays. Other related works recently studying delayed second- and high-order consensus protocols are [Gao et al., 2023; Li et al., 2024; Trindade et al., 2025].

Consider, as in [Lu and Liu, 2017, Lemma 3.1], a system of the form

$$\dot{z}_i(t) = -\sum_{j\in\mathcal{N}_i} w_{i,j}(t)\left(z_i(t) - z_j(t-\tau_{i,j}(t))\right), \tag{6}$$

where $w_{i,j}(t) \ge 0$, and each delay $\tau_{i,j}(t)$ is piecewise continuous and bounded: $\tau_{i,j}(t) \le \tau_{\max} < \infty$. Under the assumption that the time-integrated adjacency matrix

$$A(t) := \int_t^{t+T} W(\tau)\,\mathrm{d}\tau$$

induces a δ-digraph for some fixed $T > 0$ that contains a fixed root node k of a directed spanning tree for all $t \ge 0$, the system is known to reach consensus exponentially.

To express (6) compactly, we write it in functional form as

$$\dot{z}(t) = -D(t)z(t) + \mathcal{W}(z_t,t), \tag{7}$$

where $D(t)$ is a diagonal matrix and $z_t(\theta) := z(t+\theta)$ for $\theta \in [-\tau_{\max}, 0]$, following standard notation in [Hale and Lunel, 2013].

In this form, the operator $\mathcal{L}_k(z_t,t) = -D(t)z(t) + \mathcal{W}(z_t,t)$ generally fails to satisfy Assumption 1, as we show by counterexample in Appendix D. Therefore, such a delayed operator cannot be used to define $\mathcal{L}_k$ for $k < n$. However, the delayed operator may be used for the outermost function in the composition $\mathcal{L}_n$, since the

unperturbed system (6), for given delay functions, admits a unique solution. Our proof of Theorem 1 extends to this setting by interpreting ξ_n as a continuous input signal to ξ_{n-1} in (4). Provided that the remaining operators $\mathcal{L}_k$, for $k \neq n$, satisfy the assumptions of the theorem, the full solution z and its first $n-1$ derivatives are well defined and converge to consensus through the same inductive argument. A comprehensive treatment of functional differential equations can be found in [Hale and Lunel, 2013].

5. Case Studies

We now explore applications of the compositional consensus framework developed in the previous sections. In particular, we focus on second-order formation control, where each agent is modeled as a double-integrator:

$$\ddot{x} = u(x,t), \tag{8}$$

with $x(t) \in \mathbb{R}^N$. A general compositional consensus-based control law for this setting takes the form

$$u(x,t) = -\mathcal{L}_2(\dot{x} + \mathcal{L}_1(x,t),t) - \frac{\mathrm{d}}{\mathrm{d}t}\mathcal{L}_1(x,t). \tag{9}$$

This will be contrasted with a conventional second-order consensus protocol

$$u_{\mathrm{conv}}(x,t) = -\mathcal{L}_{\mathrm{vel}}(\dot{x}(t)) - \mathcal{L}_{\mathrm{pos}}(x(t)). \tag{10}$$

and what we will term a *naïve* serial consensus:

$$u_{\mathrm{ser}}(x,t) = -(\mathcal{L}_2 + \mathcal{L}_1)(\dot{x}(t)) - \mathcal{L}_2 \circ \mathcal{L}_1(x(t)), \tag{11}$$

Remark 3
In many applications, the objective is to steer a group of agents into a fixed formation and maintain a constant collective velocity. This can be achieved by introducing a desired relative position vector $d_{\mathrm{ref}} \in \mathbb{R}^N$ and a reference velocity $v_{\mathrm{ref}} \in \mathbb{R}$. To enforce the desired formation, one can work in the transformed coordinates $\tilde{x} = x - d_{\mathrm{ref}}$. This transformation does not alter the system dynamics and thus preserves the control structure. The agents will then asymptotically coordinate in the frame $\tilde{x}$, implying that $|x_i(t) - x_j(t)| \to |d_i - d_j|$ as $t \to \infty$. To ensure correct velocity tracking, a leader-follower structure may be employed, where the agents synchronize with a designated leader moving at velocity v_{ref}.

5.1 Time-varying graph Laplacians

We now consider an application where both operators in the compositional consensus protocol are defined using time-varying Laplacians: $\mathcal{L}_1(x,t) = L_1(t)x$ and $\mathcal{L}_2(x,t) = L_2(t)x$. The resulting closed-loop system becomes

$$\ddot{x} = -L_2(t)\left(\dot{x} + L_1(t)x\right) - \dot{L}_1(t)x - L_1(t)\dot{x}. \tag{12}$$

As established in Proposition 3, if $L_1(t)$ and $L_2(t)$ are piecewise continuous and sufficiently connected over time, and if $L_1(\cdot) \in C^0$, then they can be used to construct a compositional consensus protocol that guarantees second-order consensus, via Theorem 1.

Notably, such guarantees are generally not available for the corresponding naïve serial and conventional consensus protocols in (11) and (10), when the Laplacians are time-varying. We consider the following example of a string formation with time-varying connectivity.

Example 2
We define $L_{\text{path}} \in \mathbb{R}^{N\times N}$ as the Laplacian of a directed path graph with the following structure:

$$L_{\text{path}} = \begin{bmatrix} 0 & 0 & & \\ -1 & 1 & & \\ & \ddots & \ddots & \\ & & -1 & 1 \end{bmatrix}.$$

Let $L_1(t) = L_2(t) = D(t)L_{\text{path}}$, where $D(t)$ is a time-varying diagonal matrix defined as

$$[D(t)]_{i,i} = \max\{\sin(\omega_i t + \phi_i), 0\},$$

with individual frequency $\omega_i \neq 0$ and phase $\phi_i \in [0, 2\pi)$. This choice satisfies all conditions for Theorem 1 and ensures that both $L_1(t)$ and $L_2(t)$ are connected over time.

Figure 1 shows a simulation of a second-order vehicle formation with $N = 15$ agents under this protocol, with randomly chosen frequencies and phases. Despite the complexity of the system and the time-varying graph structure, the compositional protocol successfully coordinates the agents. It achieves second-order consensus, with the agents converging to their desired relative positions.

For comparison, we simulate the same formation using the time-varying versions of the conventional and naïve serial consensus controllers. The conventional controller is defined as

$$u_{\text{conv}}(x,t) = -L_1(t)\dot{x} - L_2(t)x,$$

and the result is shown in Figure 1b. The naïve serial consensus protocol is given by

$$u_{\text{ser}}(x,t) = -(L_2(t) + L_1(t))\dot{x} - L_2(t)L_1(t)x,$$

with the corresponding result shown in Figure 1c. Both alternative controllers exhibit poor transient performance: the naïve serial consensus experience an oscillatory convergence to the reference trajectory, while the conventional controller produces extreme oscillations that indicate instability.

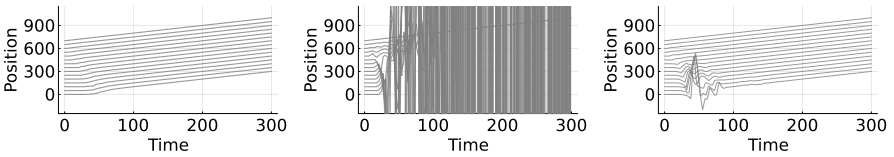

(a) Compositional consensus **(b)** Conventional consensus **(c)** Naïve serial consensus

Figure 1. Simulation results using time-varying graph Laplacians whose connection strengths vary sinusoidally. The compositional controller achieves stable second-order consensus. The naïve serial consensus exhibits a significant transient but eventually converges, while the conventional consensus seems to be truly unstable.

5.2 Saturated coordination

Time-varying consensus protocols represent just one class of systems that benefit from compositional consensus. Another important and challenging class involves nonlinear protocols, particularly those incorporating input saturation. As established in Proposition 2, the compositional consensus framework can accommodate operators such as

$$\mathcal{L}_1(x,t) = \mathrm{sat}(L_1 x), \quad \mathcal{L}_2(x,t) = \mathrm{sat}(L_2 x),$$

provided that both Laplacians contain a directed spanning tree. By contrast, no general guarantees exist for the conventional or even the naïve serial consensus protocols when such nonlinearities are present. The following example illustrates this.

Example 3
Consider formation control over a directed string network. That is, the case where $\mathcal{L}_1(x,t) = \mathcal{L}_2(x,t) = \mathrm{sat}(L_{\mathrm{path}} x)$. *The resulting control law becomes*

$$u(x,t) = -\mathrm{sat}\left(L_{\mathrm{path}}(\dot{x} + \mathrm{sat}(L_{\mathrm{path}} x))\right) - \frac{\mathrm{d}}{\mathrm{d}t}\mathrm{sat}(L_{\mathrm{path}} x).$$

The formation is initialized with a nonzero positional error to highlight the effect of saturation. The simulation results are shown in Figure 2. Despite the nonlinearities, the compositional controller (Figure 2a) achieves a smooth transition to second-order consensus.

For comparison, we simulate the same system under saturated versions of the naïve serial and conventional consensus controllers. That is

$$u_{\mathrm{ser}}(x,t) = -2\mathrm{sat}(L_{\mathrm{path}}\dot{x}) - \mathrm{sat}(L_{\mathrm{path}}\mathrm{sat}(L_{\mathrm{path}} x)),$$

and

$$u_{\mathrm{conv}}(x,t) = -\mathrm{sat}(L_{\mathrm{path}}\dot{x}) - \mathrm{sat}(L_{\mathrm{path}} x).$$

The results are shown in Figures 2c and 2b, respectively. The naïve serial consensus shows similar but slightly slower convergence than the compositional consensus.

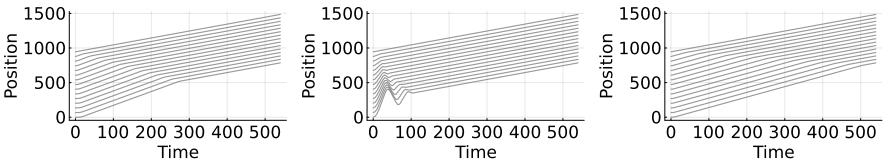

(a) Compositional consensus **(b)** Conventional consensus **(c)** Naïve serial consensus

Figure 2. Simulation of compositional, conventional, and naïve serial consensus under saturated control inputs. The compositional and naïve serial consensus protocols achieve smooth convergence, while conventional consensus exhibits an undesired transient indicative of string instability.

The saturated conventional consensus shows an indication of string instability. We have conducted larger simulations that verify this indication.

5.3 Delayed absolute feedback

We conclude with an application involving delayed absolute feedback, such as GPS-based velocity measurements. In other words, we now consider a case where Assumption 1 is relaxed. In vehicle platoons, absolute feedback has been proposed to improve performance; however, in practice, it would typically be received aperiodically and with uncertain delays. This example investigates such a scenario.

Example 4
We consider a delayed consensus protocol based on absolute measurements with static coupling weights. Written in individual-agent form, the dynamics are

$$\dot{x}_i(t) = -d_i\left(x_i(t) - x_{\mathrm{GPS}}(t-\tau_i(t))\right),$$

which can be compactly expressed as

$$\mathcal{L}_2(x_t,t) = Dx(t) - D_{\tau(t)}(\mathbf{1}x_{t,\mathrm{GPS}}),$$

where $x_t(\theta) := x(t+\theta)$ for $\theta \in [-\tau_{\max},0]$, and D is a diagonal matrix of feedback weights.

For the other operator, we use a standard linear time-invariant consensus protocol: $\mathcal{L}_1(x,t) = L_{\mathrm{path}}x(t)$. The first row of L_{path} is all zeros, modeling a virtual leader. The resulting compositional control law is

$$u_{\mathrm{comp}}(x_t,t) = -D(\dot{x} + L_{\mathrm{path}}x) + D_{\tau(t)}(\mathbf{1}\dot{x}_{t,\mathrm{GPS}}) - L_{\mathrm{path}}\dot{x}.$$

Defining $e = L_{\mathrm{path}}x$, and rearranging terms, the individual-agent control becomes

$$u_{i,\mathrm{comp}}(x_t,t) = -d_i\left(\dot{x}_i(t) - \dot{x}_{\mathrm{GPS}}(t-\tau_i(t))\right) - d_i e_i(t) - \dot{e}_i(t).$$

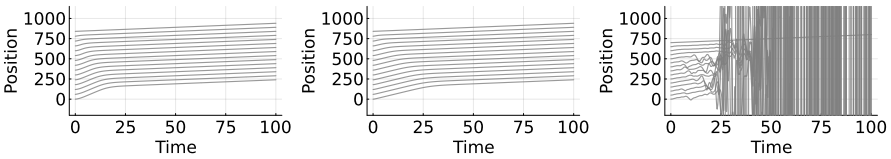

(a) Compositional consensus with delayed velocity feedback

(b) Conventional consensus with ideal velocity feedback

(c) Conventional consensus with delayed velocity feedback

Figure 3. Simulation of second-order consensus under absolute velocity feedback with and without delay. The compositional consensus controller is robust to feedback delays. This is unlike the conventional controller, which shows oscillatory behavior under the same conditions.

The last two terms correspond to standard relative feedback with local neighbors, while the first involves delayed absolute velocity feedback. This can be rewritten as

$$(\dot{x}_i(t) - \dot{x}_i(t - \tau_i(t))) - (\dot{x}_{\mathrm{GPS}}(t - \tau_i(t)) - \dot{x}_i(t - \tau_i(t)))\,,$$

which separates into two interpretable components: 1) the change in the agent's velocity since the last measurement, and 2) a delayed relative velocity signal received from the GPS.

In practice, each agent stores a record of its past velocity and periodically receives delayed GPS-based velocity references. This allows the required feedback to be implemented despite asynchronous and uncertain communication delays.

We simulate a vehicle formation with $d_i = 1$ for all agents and delays $\tau_i(t)$ sampled from a Poisson process with a mean inter-arrival time of 1 second. The compositional controller is compared to: 1) a conventional consensus protocol with perfect (non-delayed) absolute velocity feedback, and 2) the same conventional consensus protocol subject to the same delays $\tau_i(t)$ as in the compositional case.

Figure 3 shows that the compositional consensus protocol achieves smooth second-order coordination despite the delays. The conventional controller also performs well with ideal feedback, but its performance degrades significantly under delay, resulting in oscillatory behavior.

6. Conclusions

In this work, we expanded the theory of high-order coordination by introducing and analyzing a general framework for *compositional consensus*. This approach provides a flexible and modular design paradigm that accommodates practical challenges such as time-varying dynamics, nonlinearities, and communication delays. In particular, it allows us to build upon the rich literature on the convergence of first-order consensus under various non-ideal conditions and immediately apply them for higher-order

formation control. We focused on second-order coordination and vehicular formations as motivating examples, but the framework is broadly applicable. Potential applications include frequency coordination in power systems, temperature regulation in district heating networks, and large-scale multi-agent systems such as drone swarms. Implementing the controller in general presumes signaling in an n-hop neighborhood, where n is the order of the local integrator dynamics. We remark, though, that in the case $n = 2$, a nearest-neighbor implementation is also possible through a "look-ahead and look-behind protocol," see [Hansson and Tegling, 2023].

The main theoretical contribution of our work is a set of sufficient conditions for achieving asymptotic coordination in the compositional setting, formally stated in Theorem 1. These results extend the reach of classical consensus theory and offer tools for principled design in complex settings.

Future work may involve identifying necessary conditions for coordination and extending current string-stability and scalability results, which are already established for linear serial consensus, to nonlinear and time-varying compositional designs. In particular, to investigate performance guarantees that are uniform in network size. Such extensions are essential for robust and scalable deployment in large coordinated systems.

Appendix

A Proof of Lemma 1

To aid the presentation of the proof, we begin with the following supporting lemma.

Lemma 3
Let $\mathcal{L} \in C^n$ and satisfy $\mathcal{L}(z+\mathbf{1}a(t),t)$ for any integrable function $a(t)$, then, for any $k \leq n$ the following is true: $\frac{\mathrm{d}^k}{\mathrm{d}t^k}\mathcal{L}(z,t) = B_k(z,\dot{z},\dots,z^{(k)},t) = B_k(z-\mathbf{1}b_0(t),\dot{z}-\mathbf{1}b_1(t),\dots,z^{(k)}-\mathbf{1}b_k(t),t)$ for any integrable functions $b_1,b_2,\dots,b_k$.

Proof. We prove this by induction. The base case is directly proven by $B_0(z) = \mathcal{L}(z,t)$. Now, suppose it is valid for $k \leq n-1$. Firstly, taking the partial derivatives for any $j \leq k$: $\partial_{z^{(j)}} B_k(z,\dots,z^{(k)},t) = \partial_{z^{(j)}} B_k(z-\mathbf{1}b_0(t),\dots,z^{(k)}-\mathbf{1}b_k(t),t)$. This shows that all partial derivatives are invariant to arbitrary translation along the consensus, and the same argument also holds for the partial time derivative. Now, taking the time derivative of the consensus translated equation results in:

$$\frac{\mathrm{d}^{k+1}}{\mathrm{d}t^{k+1}}\mathcal{L}(z,t) = \partial_t B_k(z-\mathbf{1}c_0(t),\dots,z^{(k)}-\mathbf{1}c_k(t),t)+ \sum_{j=0}^{k} \partial_{z^{(j)}} B_k(z-\mathbf{1}c_0(t),\dots,z^{(k)}-\mathbf{1}c_k(t),t)(z^{(j+1)}-\mathbf{1}\dot{c}_j(t))$$

Let $c_j(t) = \int_0^t b_{j+1}(\tau)\mathrm{d}\tau$ and for each partial derivative term do the outlined translation to get

$$\frac{\mathrm{d}^{k+1}}{\mathrm{d}t^{k+1}}\mathcal{L}(z,t) = \partial_t B_k(z - \mathbf{1}b_0(t),\dots,z^{(k)} - \mathbf{1}b_k(t),t) + \sum_{j=0}^{k} \partial_{z^{(j)}} B_k(z - \mathbf{1}b_0(t),\dots,z^{(k)} - \mathbf{1}b_k(t),t)(z^{(j+1)} - \mathbf{1}b_{j+1}(t)).$$

This shows the sought translation invariance and concludes the proof. □

We now proceed to the proof of Lemma 1.

Proof. First, we establish that the initial condition of ξ_k is uniquely determined by the initial condition x and its first $n-1$ derivatives. By using the relation $x = \xi_1$ and (4), the following relation can be derived

$$x^{(j)} = -\frac{\mathrm{d}^{j-1}}{\mathrm{d}t^{j-1}}\mathcal{L}_1(\xi_1,t) - \cdots - \mathcal{L}_j(\xi_j,t) + \xi_{j+1}.$$

Due to the smoothness of $\mathcal{L}_k$, it is possible to expand the time derivatives in terms of the partial derivatives through the chain rule. The time derivatives can thus be expressed as

$$\frac{\mathrm{d}^j}{\mathrm{d}t^j}\mathcal{L}_k(\xi_k,t) = B_{k,j}(\xi_k,\dot{\xi}_k,\dots,\xi_k^{(j)},t).$$

Since $\dot{\xi}_k = -\mathcal{L}_k(\xi_k,t) + \xi_{k+1}$, it is possible to reduce the derivative dependence recursively and to show that

$$\frac{\mathrm{d}^j}{\mathrm{d}t^j}\mathcal{L}_k(\xi_k,t) = \hat{B}_{k,j}(\xi_k,\xi_{k+1},\dots,\xi_{k+j},t).$$

Applying this to the general case leads to

$$\xi_{j+1} = x^{(j)} + \sum_{k=1}^{j} \hat{B}_{k,j-k}(\xi_k,\dots,\xi_j,t),$$

for $j = 0,\dots,n-1$. Evaluating this at $t = 0$ shows that ξ_{j+1} is uniquely determined by $x^{(j)}(0)$ and the initial conditions of $\xi_k(0)$ for $k \leq j$. This, together with $\xi_1(0) = x(0)$, can be used to prove that $\xi_k(0)$ is uniquely determined by the initial condition of x and its derivatives through a simple induction hypothesis. An analogous proof can be made in the reverse direction and, therefore, is omitted.

For the second part of the proof, we will show that consensus of the states ξ_k implies that x achieves n^{th}-order consensus. As induction hypothesis, assume that $\|\xi_k^{(j)} - \mathbf{1}a_{k+j}(t)\| \to 0$, with the induction step taken in the j direction. The base case follows from the assumption that ξ_k all reach a consensus, that is, $\|\xi_k - \mathbf{1}a_k(t)\| \to 0$.

For the induction step consider the general expression for $\xi_k^{(j+1)}$ for $k+j \leq n-1$, which is

$$\xi_k^{(j+1)} = -\frac{\mathrm{d}^j}{\mathrm{d}t^j}\mathcal{L}_k(\xi_k,t) + \frac{\mathrm{d}^j}{\mathrm{d}t^j}\xi_{k+1}.$$

Using Lemma 3, this can be expressed in terms of the translated states

$$\xi_k^{(j+1)} = -B_{k,j}(\xi_k - \mathbf{1}a_k(t), \ldots, \xi_k^{(j)} - \mathbf{1}a_{k+j}, t) + \xi_{k+1}^{(j)}.$$

By the premise of the theorem, $\|B_{k,j}\|$ can be bounded by $\alpha_{k,j} \in \mathcal{K}$. Subtracting $\mathbf{1}a_{k+j+1}$ on both sides, taking the norm, using the triangle inequality, and bounding using the class $\mathcal{K}$ function $\alpha_{k,j}$ leads to

$$\|\xi_k^{(j+1)} - \mathbf{1}a_{k+j+1}\| \leq \|\xi_{k+1}^{(j)} - \mathbf{1}a_{k+j+1}(t)\| + \alpha_{k,j}\left(\max\left\{\|\xi_k - \mathbf{1}a_k(t)\|, \ldots, \|\xi_k^{(j)} - \mathbf{1}a_{k+j}\|\right\}\right)$$

Now, taking the limits on both sides, using the induction hypothesis together with the continuity of $\alpha_{k,j}$ shows that $\lim_{t\to\infty}\|\xi_k^{(j+1)} - \mathbf{1}a_{k+j+1}(t)\| = 0$. Thus ξ_k achieves an $(n-k+1)^{\text{th}}$-order consensus. Since $x(t) = \xi_1(t)$ it follows that x achieves an n^{th}-order consensus.

The other direction, that is, x achieving n^{th} order consensus implying that ξ_k achieves consensus is conducted similarly. Now, the induction hypothesis is that $\xi_k^{(j)} - \mathbf{1}a_{j+k}(t)$ where this will be proved by induction steps in k. First, using the relation of $x(t) = \xi_1(t)$ shows that ξ_1 achieves n^{th} order consensus. The n^{th} order consensus implies that $\|\xi^{(j)} - \mathbf{1}a_j(t)\| \to 0$ for some functions $a_j(t)$. For the induction step, we consider the relation

$$\dot{\xi}_k = -\mathcal{L}_k(\xi_k,t) + \xi_{k+1}$$

This can be differentiated $j \leq n-k-1$ times, and be rearranged to

$$\xi_{k+1}^{(j)} - \mathbf{1}a_{j+k+1} = \xi_k^{(j+1)} - \mathbf{1}a_{j+k+1} + \frac{\mathrm{d}^j}{\mathrm{d}t^j}\mathcal{L}_k(\xi_k,t).$$

Now, Lemma 3 is used to express $\frac{\mathrm{d}^j}{\mathrm{d}t^j}\mathcal{L}_k$ in terms of $B_{k,j}$ and in particular in the translated states

$$\frac{\mathrm{d}^j}{\mathrm{d}t^j}\mathcal{L}_k(\xi_k,t) = B_{k,j}(\xi_k - \mathbf{1}a_k(t), \ldots, \xi_k^{(j)}(t) - \mathbf{1}a_{k+j}(t), t).$$

Applying the triangle inequality, bounding $\|B_{k,j}\|$ with $\alpha_{k,j}$, and concluding by taking the limit shows that

$$\lim_{t\to\infty}\|\xi_{k+1}^{(j)} - \mathbf{1}a_{j+k+1}\| = 0.$$

This proves the induction step. Since this also shows that $\|\xi_k - \mathbf{1}a_k(t)\| \to 0$, we can conclude that the states will achieve consensus. □

B Proof of Lemma 2

Proof. The existence and uniqueness of ξ_n is part of the lemma premise. For the remaining states it is simple to verify that Assumption 2 implies that (4) is globally Lipschitz in ξ_k and piecewise continuous in t. Existence and uniqueness follow from a standard application of Carathéodory's existence and uniqueness theorem.

Through induction, we'll prove that consensus will be reached, that is $\|\xi_k - \mathbf{1}a_k(t)\| \to 0$ for some functions $a_k(\cdot)$. The base case with $\|\xi_n - \mathbf{1}a_n(t)\| \to 0$ follows from our assumption. Suppose it is true for all ξ_k where $k \geq p+1$. The solution for ξ_p satisfy

$$\dot{\xi}_p(t) = -\mathcal{L}_p(\xi_p(t), t) + \xi_{p+1}(t)$$

Now, subtracting the asymptotic solution of ξ_{p+1} from both sides and using the fundamental theorem of calculus results in

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\xi_p(t) - \mathbf{1}\int_0^t a_{p+1}(\tau)\mathrm{d}\tau\right) = -\mathcal{L}_p\left(\xi_p(t) - \mathbf{1}\int_0^t a_{p+1}(\tau)\mathrm{d}\tau, t\right) + \xi_{p+1} - \mathbf{1}a_{p+1}(t),$$

where the invariance of $\mathcal{L}_p$ through Assumption 1 was also used. Let $z_p(t) = \xi_p(t) - \mathbf{1}\int_0^t a_{p+1}(\tau)\mathrm{d}\tau$ and $w_p(t) = \xi_{p+1} - \mathbf{1}a_{p+1}(t)$. Then z_p satisfies

$$\dot{z}_p = -\mathcal{L}_p(z_p, t) + w_p(t)$$

where $\|w_p(t)\| \to 0$, which allows us to apply Assumption 3. In particular, there is a time T_p such that $\|w_p(t)\| \leq M_p$, where this system is ISS for some seminorm $|||\cdot|||$. To assert that $\lim_{t\to\infty} \|z_p(t) - \mathbf{1}b_p(t)\| = 0$ we can use the ε and T definition for the limit. For any $\varepsilon > 0$, choose T_p' such that $\|w_p(t)\| < \gamma_k^{-1}(\varepsilon/2)$ for all $t > T_p'$. Now, using the ISS property starting at T_p', we get

$$|||z_p(t)||| < \beta_p(|||z_p(T_p')|||, t) + \frac{\varepsilon}{2}$$

By the definition of β_p it's possible to choose a time $T \geq T_p'$ such that $\beta_p(|||z_p(T_p')|||, t) < \varepsilon/2$. This proves that the seminorm converges and in particular that there exists a $b_p(t)$ such that $\|z_p(t) - \mathbf{1}b_p(t)\| \to 0 \implies \|\xi_p(t) - \mathbf{1}\left(b_p(t) + \int_0^t a_{p+1}(\tau)\mathrm{d}\tau\right)\| \to 0$, letting $a_p = b_p(t) + \int_0^t a_{p+1}(\tau)\mathrm{d}\tau$ concludes the induction step. □

C Proof of Proposition 2

Proof. We begin by proving the result for a leader-follower network. In this case, the dynamics can be rewritten as

$$L\dot{z} = -L\,\mathrm{sat}(z) + Ld(t).$$

This representation follows from left-multiplying the system by the Laplacian L. While L is singular, we are only interested in the evolution of the disagreement vector $e = Lz$, which remains orthogonal to the consensus subspace.

A leader-follower network has a unique leader agent whose state remains unaffected by the others. For this agent, $[Lz]_i = 0$ for all time. Our goal is to show that, under sufficiently small disturbances, all other agents enter and remain in the linear regime, i.e., $|[Lz]_i(t)| < 1$ for all $t \geq T_i$.

We proceed by induction along a simple directed path of influence from the leader to any follower. Let the path consist of $m+1$ agents labeled $p_0, p_1, \ldots, p_m$, with p_0 being the leader.

Base case: The leader agent satisfies $e_{p_0}(t) = 0$ for all t, so it is trivially in the linear regime.

Inductive step: Suppose that agent p_k satisfies $|e_{p_k}(t)| < 1$ for all $t \geq T_k$. We aim to show that agent p_{k+1} enters the linear regime in finite time T_{k+1}.

The dynamics of agent p_{k+1} are given by

$$\dot{e}_{p_{k+1}} = d_{p_{k+1}} - \sum_{j \in \mathcal{N}_{p_{k+1}}} w_{p_{k+1},j} \left(\mathrm{sat}(e_{p_{k+1}}) - \mathrm{sat}(e_j)\right).$$

Since agent p_k is in the linear regime after time T_k, we have $\mathrm{sat}(e_{p_k}) = e_{p_k}$ for $t \geq T_k$. Without loss of generality, assume $e_{p_{k+1}}(T_k) > 0$ (the argument is symmetric for the negative case). We upper-bound the dynamics as

$$\dot{e}_{p_{k+1}} \leq |d_{p_{k+1}}| + w_{p_{k+1},p_k}|e_{p_k}| - w\,\mathrm{sat}(e_{p_{k+1}}) + w - w_{p_{k+1},p_k},$$

where $w = \sum_j w_{p_{k+1},j}$ is the total weight of incoming edges to agent p_{k+1}.

Now, if the disturbance is sufficiently small so that

$$|d_{p_{k+1}}| < w_{p_{k+1},p_k}(1 - |e_{p_k}|),$$

then the right-hand side of the inequality becomes negative whenever $e_{p_{k+1}} \geq 1$, implying that the agent must enter the region $|e_{p_{k+1}}| < 1$ in finite time T'_{k+1}.

After entering the linear regime, the dynamics simplify, and the state can be upper-bounded by a linear system with equilibrium state

$$e^*_{p_{k+1}} = \frac{|d_{p_{k+1}}| + w_{p_{k+1},p_k}|e_{p_k}| + w - w_{p_{k+1},p_k}}{w},$$

which can be made strictly less than 1 by choosing $d_{p_{k+1}}$ sufficiently small. The state will converge towards this bound and reach any point between this and 1 in finite time $T_{k+1} > T'_{k+1}$, and then remain there for all future time. If the agent started below this steady-state bound, the same conclusion holds with $T_{k+1} = T_k$. This completes the inductive step.

Since all agents are connected by a finite directed path originating from the leader, each agent enters the linear regime in finite time. Once all agents lie within the linear region, the local ISS result from Proposition 1 can be applied to show that the disturbance $d(t)$ has a bounded effect on $\|Lz(t)\|_\infty$.

To prove the general case, it suffices to show that the agents within the unique strongly connected component (SCC) of the graph underlying L enter and remain in the linear regime for all $t \geq T$ provided the input $d(t)$ is sufficiently small. By definition, the agents in this component evolve independently of the remaining agents.

Without loss of generality, consider the subgraph corresponding to the SCC, with Laplacian $\tilde{L} \in \mathbb{R}^{K\times K}$, where $K \geq 2$. Since the subgraph is strongly connected, its zero eigenvalue is simple, and the corresponding left Perron eigenvector w can be chosen to have strictly positive entries.

Define the diagonal matrix $W = \mathrm{diag}(w)$. Then, the matrix $\tilde{L}^\top W$ satisfies $\tilde{L}^\top W\mathbf{1} = 0$, i.e., it is a graph Laplacian of a strongly connected graph.

Consider the Lyapunov function

$$V(t) = \frac{\tilde{z}^\top \tilde{L}^\top W\tilde{z}}{2}$$

This function is non-negative and satisfies $V(t) = 0 \iff \tilde{z} \in \mathrm{span}(\mathbf{1})$, i.e., consensus.

Define $\tilde{e} = \tilde{L}\tilde{z}$. Then, the time derivative of V along trajectories of the system is

$$\dot{V} = -\tilde{z}^\top \tilde{L}^\top W\,\mathrm{sat}(\tilde{L}z) + \tilde{z}^\top \tilde{L}^\top \tilde{d}(t).$$

Expanding this, and defining $\tilde{e} = \tilde{L}\tilde{z}$, we get

$$\dot{V}(t) = \frac{\dot{V}(t)}{2} = -\sum_{i=1}^{K} w_i|\tilde{e}_i|(|\mathrm{sat}(e_i)| - \tilde{d}_i).$$

Now we seek to ensure that $\dot{V}(t) \leq -\varepsilon$ for some $\varepsilon > 0$ whenever $\|\tilde{e}\|_\infty \geq r$, for some $0 < r < 1$. To that end, note that the above can be conservatively upper bounded as

$$\dot{V} \leq -w_{\min} r(r - d_{\max}) + K w_{\max} d_{\max}^2,$$

where $w_{\min} = \min_i w_i$ and $w_{\max} = \max_i w_i$. A $d_{\max}$ that ensures that this upper bound is smaller than $-\varepsilon$ can be found as long as $w_{\min} r^2 - \varepsilon > 0$. Hence, all agents in the SCC enter the region $\|\tilde{e}\|_\infty < r$ in finite time and remain there for all future time.

Finally, since the remaining agents are influenced by at least one agent in the SCC, the same inductive argument from the leader-follower case (applied to the

directed influence paths originating from the SCC) shows that all agents eventually enter and remain in the linear regime, completing the proof. □

D Time-Delayed Consensus Protocols

Here we will illustrate the consequence of the delayed consensus protocol $\mathcal{L}_1(z,t) = Dz - \mathcal{W}(z,t)$ not satisfying Assumption 1. For simplicity, consider a two-agent system, where one is a leader and both have a constant and equal input w_0. The dynamics are then

$$\begin{aligned}\dot{z}_0 &= w_0\\ \dot{z}_1 &= -z_1(t) + z_0(t-\tau(t)) + w_0\end{aligned}$$

This system can be explicitly solved for z_0 and has the solution $z_0(t) = at + z_0(0)$. For the second, consider the case where $\tau(t) = t$ for $t \leq \tau_{\max}$. Then the solution for $t \leq \tau_{\max}$ is $z_1(t) = e^{-t}z_1(0) + (a + z_0(0))(1 - e^{-t})$. Now, provided that the system is initiated at consensus, that is $z_1(0) = z_0(0)$, then we see that the agents drift away from each other as $z_0(t) - z_1(t) = a(t - 1 + e^{-t})$. This shows that the consensus is not an equilibrium solution of this system. Therefore, we cannot expect the agents to reach a consensus when using a delayed consensus protocol for anything other than $\mathcal{L}_n$ in the compositional consensus (4).

Acknowledgements

The authors are with the Department of Automatic Control and the ELLIIT Strategic Research Area at Lund University, Lund, Sweden. This work was partially funded by Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation and the Swedish Research Council through Grant 2019-00691.

References

Aldana-López, R., R. Aragüés, and C. Sagüés (2022). "Redcho: robust exact dynamic consensus of high order". *Automatica* **141**, p. 110320.

Arcak, M. (2007). "Passivity as a design tool for group coordination". *IEEE Trans. Autom. Control* **52**:8, pp. 1380–1390.

Bamieh, B., M. R. Jovanovic, P. Mitra, and S. Patterson (2008). "Effect of topological dimension on rigidity of vehicle formations: fundamental limitations of local feedback". In: *Proc. IEEE Conf. Decis. Control*, pp. 369–374.

Bamieh, B., M. R. Jovanovic, P. Mitra, and S. Patterson (2012). "Coherence in large-scale networks: dimension-dependent limitations of local feedback". *IEEE Trans. Autom. Control* **57**:9, pp. 2235–2249.

Bullo, F. (2024). *Contraction Theory for Dynamical Systems*. 1.2. Kindle Direct Publishing.

Carnevale, G., I. Notarnicola, L. Marconi, and G. Notarstefano (2023). "Triggered gradient tracking for asynchronous distributed optimization". *Automatica* **147**, p. 110726.

Chu, K.-C. (1974a). "Decentralized control of high-speed vehicular strings". *Transp. Sci.* **8**:4, pp. 361–384.

Chu, K.-C. (1974b). "Optimal dencentralized regulation for a string of coupled systems". *IEEE Trans. Autom. Control* **19**:3, pp. 243–246.

DeGroot, M. H. (1974). "Reaching a consensus". *J. Am. Stat. Assoc.* **69**:345, pp. 118–121.

Dhullipalla, M. H. and T. Chen (2024). "A continuous-time gradient-tracking algorithm for directed networks". *IEEE Control Syst. Lett.* **8**, pp. 2199–2204.

Fax, J. A. and R. M. Murray (2004). "Information flow and cooperative control of vehicle formations". *IEEE Trans. Autom. Control* **49**:9, pp. 1465–1476.

Gao, Q., J. Cai, R. Cepeda-Gomez, and W. Xu (2023). "Improved frequency sweeping technique and stability analysis of the second-order consensus protocol with distributed delays". *Int. J. Control* **96**:2, pp. 461–474.

Hale, J. K. and S. M. V. Lunel (2013). *Introduction to functional differential equations*. Vol. 99. Springer Science & Business Media.

Hansson, J. and E. Tegling (2023). "A closed-loop design for scalable high-order consensus". In: *Proc. IEEE Conf. Decis. Control*, pp. 7388–7394.

Hansson, J. and E. Tegling (2024). "Performance bounds for multi-vehicle networks with local integrators". *IEEE Control Syst. Lett.* **8**, pp. 2901–2906.

Hansson, J. and E. Tegling (2025). "Closed-loop design for scalable performance of vehicular formations". *IEEE Trans. Control Netw. Syst.*, pp. 1–10.

He, W. and J. Cao (2011). "Consensus control for high-order multi-agent systems". *IET Control Theory Appl.* **5**, pp. 231–238.

Hinrichsen, D. and A. J. Pritchard (2005). *Mathematical systems theory I: modelling, state space analysis, stability and robustness*. Vol. 48. Springer Science & Business Media.

Jadbabaie, A., J. Lin, and A. S. Morse (2003). "Coordination of groups of mobile autonomous agents using nearest neighbor rules". *IEEE Trans. Autom. Control* **48**:6, pp. 988–1001.

Khalil, H. K. (2002). *Nonlinear Systems*. 3rd. Prentice Hall, Upper Saddle River, NJ.

Kia, S. S., B. Van Scoy, J. Cortes, R. A. Freeman, K. M. Lynch, and S. Martinez (2019). "Tutorial on dynamic average consensus: the problem, its applications, and the algorithms". *IEEE Control Syst. Mag.* **39**:3, pp. 40–72.

Levine, W. and M. Athans (1966). "On the optimal error regulation of a string of moving vehicles". *IEEE Trans. Autom. Control* **11**:3, pp. 355–361.

Li, C.-J., G.-P. Liu, P. He, F. Deng, and J. Cao (2024). "Consensus of multiple high-order integrator agents with time-varying connectivity and delays: protocols using only relative states". *IEEE Trans. Syst. Man Cybern.: Syst.* **54**:5, pp. 2663–2675.

Li, C.-J., G.-P. Liu, P. He, F. Deng, and H. Li (2022). "Dynamic consensus of second-order networked multiagent systems with switching topology and time-varying delays". *IEEE Trans. Cybern.* **52**:11, pp. 11747–11757.

Li, Y., J. Xiang, and W. Wei (2011). "Consensus problems for linear time-invariant multi-agent systems with saturation constraints". *IET Control Theory Appl.* **5**, pp. 823–829.

Lin, Z., B. Francis, and M. Maggiore (2005). "Necessary and sufficient graphical conditions for formation control of unicycles". *IEEE Trans. Autom. Control* **50**:1, pp. 121–127.

Liu, Y. and Y. Jia (2010). "Consensus problem of high-order multi-agent systems with external disturbances: an H-infinity analysis approach". *Int. J. Robust Nonlinear Control* **20**:14, pp. 1579–1593.

Lu, M. and L. Liu (2017). "Distributed feedforward approach to cooperative output regulation subject to communication delays and switching networks". *IEEE Trans. Autom. Control* **62**:4, pp. 1999–2005.

Lyu, J., J. Qin, D. Gao, and Q. Liu (2016). "Consensus for constrained multi-agent systems with input saturation". *Int. J. Robust Nonlinear Control* **26**:14, pp. 2977–2993.

Melzer, S. and B. Kuo (1971). "Optimal regulation of systems described by a countably infinite number of objects". *Automatica* **7**:3, pp. 359–366.

Moreau, L. (2005). "Stability of multiagent systems with time-dependent communication links". *IEEE Trans. Autom. Control* **50**:2, pp. 169–182.

Moreau, L. (2004). "Stability of continuous-time distributed consensus algorithms". In: *Proc. IEEE Conf. Decis. Control*. Vol. 4, pp. 3998–4003.

Olfati-Saber, R. (2006). "Flocking for multi-agent dynamic systems: algorithms and theory". *IEEE Trans. Autom. Control* **51**:3, pp. 401–420.

Olfati-Saber, R. and R. Murray (2004). "Consensus problems in networks of agents with switching topology and time-delays". *IEEE Trans. Autom. Control* **49**:9, pp. 1520–1533.

Olfati-Saber, R., J. A. Fax, and R. M. Murray (2007). "Consensus and cooperation in networked multi-agent systems". *Proc. IEEE* **95**:1, pp. 215–233.

Ren, W. and R. Beard (2005). "Consensus seeking in multiagent systems under dynamically changing interaction topologies". *IEEE Trans. Autom. Control* **50**:5, pp. 655–661.

Ren, W., R. W. Beard, and T. W. McLain (2005). "Coordination variables and consensus building in multiple vehicle systems". In: Kumar, V. et al. (Eds.). *Cooperative Control: A Post-Workshop Volume 2003 Block Island Workshop on Cooperative Control*. Springer, Berlin, pp. 171–188.

Ren, W., K. Moore, and Y. Q. Chen (2006). "High-order consensus algorithms in cooperative vehicle systems". In: *Proc. IEEE Int. Conf. Netw. Sens. Control*, pp. 457–462.

Ren, W., K. L. Moore, and Y. Q. Chen (2007). "High-order and model reference consensus algorithms in cooperative control of multi-vehicle systems". *J. Dyn. Syst. Meas. Control* **129**:5, pp. 678–688.

Rezaee, H. and F. Abdollahi (2015). "Average consensus over high-order multiagent systems". *IEEE Trans. Autom. Control* **60**:11, pp. 3047–3052.

Spanos, D. P., R. Olfati-Saber, and R. M. Murray (2005). "Dynamic consensus on mobile networks". In: *Proc. IFAC World Congress*, pp. 1–6.

Stüdli, S., M. Seron, and R. Middleton (2017a). "From vehicular platoons to general networked systems: string stability and related concepts". *Annu. Rev. Control* **44**, pp. 157–172.

Stüdli, S., M. M. Seron, and R. H. Middleton (2017b). "Vehicular platoons in cyclic interconnections with constant inter-vehicle spacing". In: *Proc. IFAC World congress*, pp. 2511–2516.

Tegling, E., B. Bamieh, and H. Sandberg (2023). "Scale fragilities in localized consensus dynamics". *Automatica* **153**, p. 111046.

Tegling, E., R. H. Middleton, and M. M. Seron (2019). "Scalability and fragility in bounded-degree consensus networks". *IFAC-PapersOnLine* **52**:20, pp. 85–90.

Tian, Y.-P. and Y. Zhang (2012). "High-order consensus of heterogeneous multi-agent systems with unknown communication delays". *Automatica* **48**:6, pp. 1205–1212.

Trindade, P., P. Batista, and R. Cunha (2025). "Delay-margin design in consensus for high-order integrator agents". *Automatica* **174**, p. 112173.

Trindade, P., R. Cunha, and P. Batista (2024). "Optimal consensus for high-order integrator agents". In: *Proc. IEEE Conf. Decis. Control*, pp. 6957–6962.

Yadlapalli, S. K., S. Darbha, and K. R. Rajagopal (2006). "Information flow and its relation to stability of the motion of vehicles in a rigid formation". *IEEE Trans. Autom. Control* **51**:8, pp. 1315–1319.

Paper V

Input-Output Pseudospectral Bounds for Transient Analysis of Networked and High-Order Systems

Jonas Hansson **Emma Tegling**

Abstract

Motivated by a need to characterize transient behaviors in large network systems in terms of relevant signal norms and worst-case input scenarios, we propose a novel approach based on existing theory for matrix pseudospectra. We extend pseudospectral theorems, pertaining to matrix exponentials, to an input-output setting, where matrix exponentials are pre- and post-multiplied by input and output matrices. Analyzing the resulting transfer functions in the complex plane allows us to state new upper and lower bounds on system transients. These are useful for higher-order matrix differential equations, and specifically control of double-integrator networks such as vehicle formation problems. Therefore, we illustrate the theory's applicability to the problem of vehicle platooning and the question of string stability, and show how unfavorable transient behaviors can be discerned and quantified directly from the input-output pseudospectra.

1. Introduction

Characterizing dynamic properties of systems with structure, in particular, network structure, is a long-standing problem in the field. While questions of stability and convergence have dominated the literature since the early works [Fax and Murray, 2004; Olfati-Saber and Murray, 2004], important questions pertaining to the performance and robustness of network systems are increasingly gaining attention. For example, [Bamieh et al., 2012] and later [Siami and Motee, 2016; Tegling et al., 2019] have described fundamental limitations to the performance of large networks subject to structural (sparsity) constraints, stated in terms of system norms.

A particular area where dynamic behaviors have received more attention is that of vehicle platooning, that is, the control of strings of vehicles, see [Levine and Athans, 1966; Chu, 1974] for early works. Here, it is fundamentally important to prevent disturbance propagation through the string (to avoid collisions!), and therefore, to have uniform bounds on error amplifications during transients. This has motivated the notion of *string stability*, see e.g., [Swaroop and Hedrick, 1996; Seiler et al., 2004] or [Stüdli et al., 2017; Feng et al., 2019] for more recent surveys. Conditions for string stability fall, roughly speaking, into two categories: 1) bounding the amplification of a disturbance from vehicle i to vehicle j, or 2) requiring that bounded initial errors lead to bounded output errors, independently of the string length. The choice of signal norms, however, is central for the bounds in this literature, and the interpretations they allow for. Many works have done analyses based on $\mathcal{L}_2$ to $\mathcal{L}_2$ string stability, see [Herman et al., 2017; Seiler et al., 2004; Ploeg et al., 2014] while the, as argued e.g. in [Feintuch and Francis, 2012], possibly more important $\mathcal{L}_\infty$ to $\mathcal{L}_\infty$ disturbance amplification has received significantly less attention even if considered in [Swaroop and Hedrick, 1996; Chu, 1974]. In this work, we shed light on a new approach to analyzing such bounds for input-output systems in general, and networks and vehicle strings in particular.

This approach takes off from the literature on *pseudospectra*. Pseudospectra, which complement spectral analysis of linear systems, especially for those with non-normal operators, have seen usage in describing the transient behavior of both differential and difference equations. The works are too numerous to mention, but we refer to [Trefethen and Embree, 2005] for an excellent textbook on the subject. Through pseudospectra one can state lower and upper bounds on the transient of the exponential matrix, i.e., on $\sup_{t\geq 0}\|e^{t\mathcal{A}}\|$, and thereby on the solution to a linear differential equation. In other words, on the transient response of the internal states of a linear system. The most famous such bounds are given by the Kreiss theorem [Kreiss, 1962]. However, in control, and in particular, network applications including vehicle platooning, we are not necessarily interested in the transients of the internal states. For instance, vehicular formation dynamics tend to have a double integrator rendering certain internal states unbounded, while inter-vehicular distances may be well-behaved. To cope with this one can incorporate measurement and input matrices $\mathcal{C},\mathcal{B}$ and then bound $\sup_{t\geq 0}\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\|$ instead.

The extension of pseudospectral bounds to such an input-output setting is the main focus of the present work. For this purpose we will define a notion of *input-output pseudospectra*. These will, in the case of higher-order systems (by which we mean systems with more than one integrator), become closely related to *structured* pseudospectra, which have been studied in [Tisseur and Higham, 2001; Lancaster and Psarrakos, 2005] and applied to mechanical systems in [Green et al., 2006]. In these works the main focus has been on the robustness of solutions to matrix polynomial equations including the quadratic eigenvalue problem. The related analysis of transient behavior of $\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\|$ has, to the best of our knowledge, barely received attention, though some structured Kreiss-like theorems were proven in [Matsuo, 1994; Plischke, 2005].

This paper aims to highlight the potential usefulness of the pseudospectral framework for networked systems and systems with higher-order dynamics. Platooning, where vehicles are modeled as double integrators (the acceleration is actuated), and which have a string network topology, is a prototypical example. We first generalize certain key results from [Trefethen and Embree, 2005] to an input-output setting. Furthermore, we use complex analysis to derive new upper bounds on the transients of state space realizations, which are especially useful for systems that have high-order dynamics. The generalizations lead to lower and upper bounds on the transient $\sup_{t\geq 0}\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\|$, which under given input scenarios imply bounds on the output $\sup_{t\geq 0}\|y(t)\|$ (in any p-norm). Through examples we show how the new bounds can be applied. For a large-scale platooning problem, we compute bounds on the deviations from equilibrium for a worst-case bounded initial condition.

The remainder of this paper is organized as follows. In Section 2 we introduce the preliminaries of this work. Lower and upper bounds on the transient of $\sup_t\|y(t)\|$ and simple examples illustrating how to apply the bounds are presented in Section 3. Then we illustrate an application of our results in the form of vehicle strings in Section 4. Lastly our conlusions are presented in Section 5.

2. Preliminaries

Consider the linear time-invariant system

$$\begin{aligned}\dot{\xi}(t) &= \mathcal{A}\xi(t) + \mathcal{B}u(t)\\ y(t) &= \mathcal{C}\xi(t),\end{aligned} \tag{1}$$

where the state $\xi \in \mathbb{R}^N$, $\mathcal{A} \in \mathbb{R}^{N\times N}$, $\mathcal{B} \in \mathbb{R}^{N\times P}$, $\mathcal{C} \in \mathbb{R}^{Q\times N}$, and output $y \in \mathbb{R}^Q$. The initial condition is $\xi(0) = \xi_0$. We will interpret $\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}$ as a transfer matrix and call the system (1) input-output stable if all poles of this transfer matrix lie in the open left half plane. Denote by $\sigma(\mathcal{A})$ the spectrum, i.e., the set of eigenvalues of $\mathcal{A}$.

We will often let the system in (1) model matrix differential equations of the form

$$\begin{aligned} x^{(l)}(t) + A_{l-1}x^{(l-1)}(t) + \cdots + A_0 x(t) &= Bu(t) \\ y(t) &= C\xi(t), \end{aligned} \tag{2}$$

where $x(t) \in \mathbb{R}^n$ and $x^{(k)}$ denotes the k^{th} time derivative of x: $x^{(k)}(t) = \frac{\mathrm{d}^k x(t)}{\mathrm{d}t^k}$. In this case, $\xi(t) = [x, \dot{x}, \ldots, x^{(l-1)}]^\top \in \mathbb{R}^{nl}$, with $nl = N$. This system can be equivalently stated on block-companion form as

$$\begin{aligned} \dot{\xi}(t) &= \underbrace{\begin{bmatrix} 0 & I_n & 0 & \dots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & I_n \\ -A_0 & -A_1 & \dots & -A_{l-1} \end{bmatrix}}_{\mathcal{A}} \xi(t) + \underbrace{\begin{bmatrix} 0 \\ \vdots \\ 0 \\ B \end{bmatrix}}_{\mathcal{B}} u(t) \\ y(t) &= C\xi(t). \end{aligned} \tag{3}$$

2.1 Signal and system norms

Norms are central to this work. Here we will consider the standard vector p-norms:

$$\|x\|_p = \begin{cases} \left(\sum_{k=1}^N |x_k|^p\right)^{\frac{1}{p}} & \text{if } 1 \le p < \infty \\ \max_k |x_k| & \text{if } p = \infty, \end{cases}$$

where $x \in \mathbb{C}^N$. For matrices we consider the corresponding induced norms, i.e.

$$\|\mathcal{A}\| = \sup_{\|x\|=1} \|\mathcal{A}x\|,$$

where $\mathcal{A} \in \mathbb{C}^{M \times N}$.

In general, our results can be interpreted in any of these norms and we will often omit the subscript to indicate that the results are valid for all of them. What we need for our theorems is, more specifically, that the matrix norms are submultiplicative, which means that the following inequality is valid for any two compatible matrices A_1, A_2

$$\|A_1 A_2\| \le \|A_1\| \|A_2\|.$$

It is well known that this is true for all the p-norms.

2.2 Input-output scenarios

We will present bounds in terms of the scaled exponential matrix $Ce^{\mathcal{A}t}\mathcal{B}$. Its norm can be seen as bounds on the transient response of the system (1) in the following scenarios:

Impulse response Consider the input signal $\{u(t) = \delta(t)u_0\}$ with $u_0 \in \mathbb{R}^P$ and let $||u_0|| = 1$ in some norm. The solution of (1) is given by

$$y(t) = \mathcal{C}e^{t\mathcal{A}}\mathcal{B}u_0 \tag{4}$$

and the worst possible transient of $y(t)$ is given by

$$\sup_t \|y(t)\| = \sup_t \sup_{\|u_0\|=1} \|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}u_0\| = \sup_t \|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\|.$$

Response to an initial condition An initial condition response is given by

$$y(t) = \mathcal{C}e^{t\mathcal{A}}\xi(0).$$

To study the worst possible initial condition with respect to resulting deviations in the output $y(t)$ we may consider

$$\sup_t \|y(t)\| = \sup_t \sup_{\|\xi_0\|=1} \|\mathcal{C}e^{t\mathcal{A}}\xi_0\| = \sup_t \|\mathcal{C}e^{t\mathcal{A}}\|.$$

The corresponding analysis for the worst-case *structured* initial condition is done by multiplying ξ_0 by $\mathcal{B}$. In this case,

$$\sup_t \|y(t)\| = \sup_t \sup_{\|\xi_0\|=1} \|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\xi_0\| = \sup_t \|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\|.$$

For example, $\mathcal{B} = (I, 0, \ldots, 0)^T$ in (3) corresponds to all initial derivatives being zero.

2.3 Complex analysis

The basis for our upcoming theorems is three Laplace transform results, which were also used to derive key results in [Trefethen and Embree, 2005]. For completeness they are also presented.

Lemma 1— [Trefethen and Embree, 2005, Theorem 15.1]
Let $\mathcal{A}$ be a matrix. There exist $\omega \in \mathbb{R}$ and $M \geq 1$ such that

$$\|e^{t\mathcal{A}}\| \leq Me^{\omega t} \qquad \forall t \geq 0. \tag{5}$$

Any $s \in \mathbb{C}$ with $Res > \omega$ is in the resolvent set of $\mathcal{A}$, with

$$(sI - \mathcal{A})^{-1} = \int_0^\infty e^{-st}e^{t\mathcal{A}}\mathrm{d}t. \tag{6}$$

If $\mathcal{A}$ is a matrix or bounded operator, then

$$e^{t\mathcal{A}} = \frac{1}{2\pi i}\int_\Gamma e^{st}(sI - \mathcal{A})^{-1}\mathrm{d}s, \tag{7}$$

where Γ is any closed and positively oriented contour that encloses $\sigma(\mathcal{A})$ once in its interior.

2.4 Pseudospectra

Pseudospectra have proven themselves to be a useful tool for analysing the transient behavior and robustness of differential equations, see e.g. [Green et al., 2006]. There are several equivalent definitions of the pseudospectra of a matrix $\mathcal{A} \in \mathbb{C}^{N\times N}$. Two equivalent and well known are:

Definition 1—ε-pseudospectra

$$\sigma_\varepsilon(\mathcal{A}) = \{s \in \mathbb{C} \mid \|(sI-\mathcal{A})^{-1}\| > \varepsilon^{-1}\} \tag{8}$$

and

Definition 2—ε-pseudospectra

$$\sigma_\varepsilon(\mathcal{A}) = \{s \in \mathbb{C} \mid s \in \sigma(\mathcal{A}+E) \textit{ for some } E \in \mathbb{C}^{N\times N} \textit{ with } \|E\| < \varepsilon\}, \tag{9}$$

where $\sigma(\mathcal{A})$ denotes the (usual) spectrum of a matrix $\mathcal{A}$. We will also make use of the ε-pseudospectral abscissa, defined as $\alpha_\varepsilon = \sup_{s\in\sigma_\varepsilon} \mathrm{Re}s$.

From the two definitions of σ_ε we can get an idea of what they are used for. The first relates to the size of the resolvent and enables complex analysis in line with Lemma 1. The latter relates to the robustness of the matrix under perturbations. By considering level curves of pseudospectra for various ε-levels it is possible to get an understanding of the solutions of the linear differential equation $\dot{x}(t) = \mathcal{A}x(t)$ and of how sensitive the system is to perturbations.

When one is concerned with the transient behaviour of an input-output system as defined in (1) it will be proven useful to generalize Definition 1 in the following way:

Definition 3—Input-output ε-pseudospectra

$$\sigma_\varepsilon(\mathcal{A},\mathcal{B},\mathcal{C}) = \{s \in \mathbb{C} \mid \|\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}\| > \varepsilon^{-1}\}. \tag{10}$$

The corresponding input-output *pseudospectral abscissa* we define as

$$\alpha_\varepsilon(\mathcal{A},\mathcal{B},\mathcal{C}) = \sup_{s\in\sigma_\varepsilon(\mathcal{A},\mathcal{B},\mathcal{C})} \mathrm{Re}(s).$$

We also define the input-output spectrum $\sigma(\mathcal{A},\mathcal{B},\mathcal{C})$ as the set of poles of the transfer matrix $\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}$.

2.5 Kreiss theorem

The transient behavior of a matrix exponential for a stable matrix $\mathcal{A} \in \mathbb{R}^{N\times N}$ can be bounded through the so called *Kreiss bounds* [Trefethen and Embree, 2005, Thrm. 18.5]:

$$\mathcal{K}(\mathcal{A}) \leq \sup_{t\geq 0} \|e^{t\mathcal{A}}\| \leq eN\mathcal{K}(\mathcal{A}). \tag{11}$$

Here the Kreiss constant is defined as

$$\mathcal{K}(\mathcal{A}) = \sup_{\mathrm{Re}s>0} \mathrm{Re}s\|(sI-\mathcal{A})^{-1}\|. \tag{12}$$

Comparing (12) to Definition 1, the relation between the Kreiss bound and pseudospectra becomes evident. In fact, it holds that $\mathcal{K}(\mathcal{A}) = \sup_{\varepsilon>0} \alpha_\varepsilon/\varepsilon$. In a controls context, it is natural to not only consider the matrix exponential, but rather an input-output setting. We therefore define the input-output Kreiss constant

$$\mathcal{K}(\mathcal{A},\mathcal{B},\mathcal{C}) = \sup_{\mathrm{Re}s>0} \mathrm{Re}s\|\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}\| = \sup_{\varepsilon>0} \frac{\alpha_\varepsilon(\mathcal{A},\mathcal{B},\mathcal{C})}{\varepsilon}. \tag{13}$$

3. Input-Output Transient Bounds

We now make use of the theory in the previous section to derive bounds on the transient performance of the system (1), under the input-output scenarios introduced earlier. We will give both lower and upper bounds. As a starting point, consider the following proposition, which is a simple but important extension to Lemma 1:

Proposition 1
Let $\mathcal{A}$, $\mathcal{B}$ and $\mathcal{C}$ be matrices and let $\|\cdot\|$ denote a submultiplicative norm. There exist $w \in \mathbb{R}$ and $M \geq \|\mathcal{C}\mathcal{B}\|$ such that

$$\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| \leq Me^{\omega t} \qquad \forall t \geq 0. \tag{14}$$

Any $s \in \mathbb{C}$ with $Res > \omega$ is in the resolvent set of $\mathcal{A}$, with

$$\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B} = \int_0^\infty e^{-st}\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\mathrm{d}t, \tag{15}$$

$$\mathcal{C}e^{t\mathcal{A}}\mathcal{B} = \frac{1}{2\pi i}\int_\Gamma e^{st}\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}\mathrm{d}s, \tag{16}$$

and where Γ is any closed and positively oriented contour that encloses $\sigma(\mathcal{A},\mathcal{B},\mathcal{C})$ once in its interior.

Proof. First, (14) follows from the norm's submultiplicativity and (5) as

$$\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| \leq \|\mathcal{C}\|\|\mathcal{B}\|\|e^{t\mathcal{A}}\| \leq \|\mathcal{C}\|\|\mathcal{B}\|\hat{M}e^{wt},$$

with $M = \|\mathcal{C}\|\|\mathcal{B}\|\hat{M}$. Letting $t = 0$ yields $\|\mathcal{CB}\| \leq M$.

Next, (15) follows from linearity of the integral, i.e., the fact that for any compatible matrices B, C, and $f(x)$ we have $B\int(f(x)\mathrm{d}x)C = \int Bf(x)C\mathrm{d}x$.

Last, consider (16). Through linearity and (7), we get

$$\mathcal{C}e^{t\mathcal{A}}\mathcal{B} = \frac{1}{2\pi i}\int_{\Gamma'} e^{st}\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}\mathrm{d}s,$$

where Γ' encircles $\sigma(A)$. If $\sigma(\mathcal{A},\mathcal{B},\mathcal{C}) = \sigma(\mathcal{A})$ we are done. If not, suppose that there are n_p distinct poles $s_p \in \sigma(A)$ such that $s_p \notin \sigma(\mathcal{A},\mathcal{B},\mathcal{C})$. Let Γ' be the union of Γ and n_p disjoint circles with radius ε with the poles s_p at the center. Let ε be sufficiently small such that the ε-circles are disjoint from $\sigma(\mathcal{A},\mathcal{B},\mathcal{C})$. Now, since the transfer matrix $\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}$ does not contain any poles in the interior of the ε-discs, each of the transfer functions is holomorphic in each disc enclosed by the ε-circles. By the maximum modulus principle they cannot have any strict local maximum in the interior of each ε-disc. This implies that there is an $M_\varepsilon \geq 0$ such that each transfer function $|(\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B})_{i,j}| \leq M_\varepsilon$. In turn, this implies that $\|e^{st}(\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B})\|_\infty \leq e^{Re(s_p+\varepsilon)t}M_\varepsilon P$ on any circle γ, where P is the number of columns of $\mathcal{B}$. The curve integral is thus bounded by $\int_\gamma \|\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}\|_\infty \mathrm{d}s \leq Pe^{Re(s_p+\varepsilon)t}M_\varepsilon\varepsilon 2\pi$, which then converges to 0 as $\varepsilon \to 0$. This is true for all n_p circles and so we can ignore the part encircling the non-observable poles. By equivalence of norms, this is true for any p-norm. □

We will now make use of Proposition 1 to state upper and lower bounds on the quantity $\sup_{t\geq 0}\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\|$.

3.1 Lower bound

We begin by stating a lower bound analogous to the lower bound in the Kreiss theorem (11). Despite its relevance to control systems, this extension of the Kreiss theorem has, to our knowledge, not been observed in the literature apart from [Matsuo, 1994] and [Plischke, 2005]. The short proof we present here, however, is new.

Theorem 1—Lower bound

$$\sup_{t\geq 0}\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| \geq \sup_{Res>0} Res\|\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}\| \tag{17}$$

Proof. Let $M = \sup_{t\geq 0}\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\|$ and $\mathrm{Re}s > 0$. From (15) we have

$$\|\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}\| = \|\int_0^\infty e^{-st}\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\mathrm{d}t\|$$
$$\implies \|\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}\| \leq M\int_0^\infty e^{-t\mathrm{Re}s}\mathrm{d}t = \frac{M}{\mathrm{Re}s},$$

multiplying both sides by $\mathrm{Re}s$ proves the inequality. □

The theorem reveals that the input-output Kreiss constant $\mathcal{K}(\mathcal{A},\mathcal{B},\mathcal{C})$ defined in (13) can lower bound the transient of the system (1) under the input scenarios in Section 2.2.

3.2 Upper bounds

Now we present three ways to bound the transient from above, again, using Proposition 1 as a basis for the proofs.

Theorem 2—First upper bound
If $\mathcal{A}$, $\mathcal{B}$, $\mathcal{C}$ are matrices and L_ε is the arc length of the boundary of $\sigma_\varepsilon(\mathcal{A},\mathcal{B},\mathcal{C})$ or of its convex hull for some $\varepsilon > 0$, then

$$\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| \leq \frac{L_\varepsilon e^{t\alpha_\varepsilon(\mathcal{A},\mathcal{B},\mathcal{C})}}{2\pi\varepsilon}, \tag{18}$$

where $\alpha_\varepsilon(\mathcal{A},\mathcal{B},\mathcal{C}) = \sup\{Res \mid \|\mathcal{C}(s-\mathcal{A})^{-1}\mathcal{B}\| > \varepsilon^{-1}\}$.

Proof. For any closed contour Γ enclosing $\sigma_\varepsilon(\mathcal{A},\mathcal{B},\mathcal{C})$ we have (16). Taking the norm on both sides gives

$$\begin{aligned}\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| &= \left\| \frac{1}{2\pi i}\int_\Gamma e^{st}\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}\mathrm{d}s \right\| \\ &\leq \frac{1}{2\pi}\int_\Gamma \|e^{st}\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}\|\mathrm{d}s \leq \frac{L_\varepsilon e^{t\alpha_\varepsilon(\mathcal{A},\mathcal{B},\mathcal{C})}}{2\pi\varepsilon}.\end{aligned}$$

The second inequality follows since $\|\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}\| \leq \varepsilon^{-1}$ along Γ. The convex hull can be used to reduce the length L_ε of Γ. This is possible as $\|\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}\| \leq \varepsilon^{-1}$ on the boundary of the convex hull. □

Theorem 2 is a fairly straightforward extension of [Trefethen and Embree, 2005, Theorem 15.2]. However, we next present a novel alternative characterization which will prove useful, in particular for classes of higher-order matrix differential equations.

Theorem 3—Second upper bound
Let the system (1) with $(\mathcal{A},\mathcal{B},\mathcal{C})$ be input-output stable and let $R = a\|\mathcal{A}\|$ for some $a > 1$. Then

$$\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| \leq \frac{1}{2\pi}\int_{-R}^{R}\|\mathcal{C}(i\omega-\mathcal{A})^{-1}\mathcal{B}\|d\omega + \frac{\|\mathcal{C}\|\|\mathcal{B}\|}{2-2a^{-1}} \tag{19}$$

Proof. By definition of input-output stability all poles of the transfer matrix $\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}$ lie in the left half plane. Furthermore, the spectrum $\sigma(\mathcal{A})$ is contained in the disc $|s| \leq \|\mathcal{A}\|$ since for any eigenvector x of $\mathcal{A}$ with $\|x\| = 1$ we have $\|\mathcal{A}\| \geq \|\mathcal{A}x\| = |\lambda|$. Now take Γ to be the semicircle with radius $R > \|A\|$ that goes up the imaginary axis and then extends into the left half plane. Then this Γ encloses the input-output spectrum $\sigma(\mathcal{A},\mathcal{B},\mathcal{C})$ and (16) yields

$$\begin{aligned}\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| &\leq \frac{1}{2\pi}\int_{\Gamma} e^{t\mathrm{Re}(s)}\|\mathcal{C}(s-\mathcal{A})^{-1}\mathcal{B}\|\mathrm{d}s \\ &\leq \frac{1}{2\pi}\int_{-R}^{R}\|\mathcal{C}(i\omega-\mathcal{A})^{-1}\mathcal{B}\|\mathrm{d}\omega \\ &+ \frac{1}{2\pi}\int_{\pi/2}^{3\pi/2}\|\mathcal{C}(Re^{i\theta}-\mathcal{A})^{-1}\mathcal{B}\|R\mathrm{d}\theta.\end{aligned}$$

If $|s| = \|\mathcal{A}\|a$ and $a > 1$, then, the integral in the second term can be bounded using the following series expansion of the inverse:

$$\|(sI-\mathcal{A})^{-1}\| = \left\|\frac{1}{s}\sum_{k=0}^{\infty}\left(\frac{\mathcal{A}}{s}\right)^{k}\right\| \leq \frac{1}{\|\mathcal{A}\|}\frac{1}{a-1}$$

This, together with submultiplicativity yield $\|\mathcal{C}(Re^{i\theta}-\mathcal{A})^{-1}\mathcal{B}\|R \leq \|\mathcal{C}\|\|\mathcal{B}\|\frac{\|\mathcal{A}\|}{\|\mathcal{A}\|}\frac{a}{a-1}$ which can be used to upper bound the second integral to $\frac{\|\mathcal{C}\|\|\mathcal{B}\|}{2-2a^{-1}}$. □

Now we will look into another bound, similar in its nature.

Theorem 4—Third Upper bound

Let the system (1) with $(\mathcal{A},\mathcal{B},\mathcal{C})$ be input-output stable. If $\|\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}\| \leq M|s|^{-\beta}$ for all $|s| \geq K$ for some $\beta > 1$, $M > 0$, and $K > 0$. Then

$$\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| \leq \frac{1}{2\pi}\int_{-\infty}^{\infty}\|\mathcal{C}(i\omega-\mathcal{A})^{-1}\mathcal{B}\|\mathrm{d}\omega < \infty. \tag{20}$$

Proof. Taking the norm of (16), we get

$$\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| \leq \frac{1}{2\pi}\int_{\Gamma}\|e^{st}\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}\|\mathrm{d}s.$$

Now, use the same semicircle Γ as in the proof of Theorem 3 with radius $R > \|\mathcal{A}\|$. This Γ encloses the input-output spectra $\sigma(\mathcal{A},\mathcal{B},\mathcal{C})$. Furthermore if $R > K$ we have

$$\begin{aligned}\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| &\leq \frac{1}{2\pi}\lim_{R\to\infty}\left(\int_{-R}^{R}\|e^{i\omega t}\mathcal{C}(i\omega-\mathcal{A})^{-1}\mathcal{B}\|d\omega + \pi RMR^{-\beta}\right) \\ &= \frac{1}{2\pi}\int_{-\infty}^{\infty}\|\mathcal{C}(i\omega-\mathcal{A})^{-1}\mathcal{B}\|d\omega,\end{aligned}$$

where the last equality follows from the condition $\beta > 1$. □

The condition on β in Theorem 4 can be related to the relative degree of the system. For instance, if $\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B} = (s^2I + sA_1 + A_0)^{-1}$ and is input-output stable, then $\beta = 2$ and it is possible to apply the theorem.

The usefulness of the three upper bounds boils down to the fact that the spectrum is usually difficult to characterize. For the first bound (18), a good description of the pseudospectra is needed, while in the second and third bounds (19)–(20) good knowledge of the resolvent along the imaginary axis is needed. We will clarify through two simple examples.

Example 1
Consider the dynamical system

$$\begin{aligned}\dot{\xi} &= \begin{bmatrix} 0 & 1 \\ -1 & -2 \end{bmatrix} \xi + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u \\ y &= \begin{bmatrix} 1 & 0 \end{bmatrix} \xi.\end{aligned} \tag{21}$$

Suppose we are interested in the impulse response of the system. Then we have

$$\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B} = \frac{1}{(s+1)^2}.$$

In this case we can see that the ε-level curves of $\|\mathcal{C}(sI-\mathcal{A})^{-1}\mathcal{B}\| = 1/\varepsilon$ are given by the circles $|s+1| = \sqrt{\varepsilon}$. From (18) we see that the upper bound for each ε is

$$\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| \leq \frac{2\pi\sqrt{\varepsilon}e^{t(-1+\sqrt{\varepsilon})}}{2\pi\varepsilon} = \frac{e^{t(-1+\sqrt{\varepsilon})}}{\sqrt{\varepsilon}}.$$

The lowest upper bound is achieved for $\varepsilon = 1$ and is simply $\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| \leq 1$.

The third upper bound (Theorem 4) requires input-output stability, which is clearly satisfied. The relative degree is 2 which implies $\beta = 2 > 1$. To calculate the bound (20) we need to calculate the integral along the imaginary axis. In this case

$$\begin{aligned}\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| &\leq \frac{1}{2\pi}\int_{-\infty}^{\infty} \|\mathcal{C}(\omega i I - \mathcal{A})^{-1}\mathcal{B}\| \mathrm{d}\omega \\ &= \frac{1}{\pi}\int_0^{\infty} \frac{1}{\omega^2+1}\mathrm{d}\omega = \frac{1}{2}\end{aligned}$$

A lower bound of this system can be calculated by only considering the real axis (and in this case this is also optimal). This leads to optimizing

$$\sup_{t\geq 0} \|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| \geq \sup_{x>0} \left\| \frac{x}{(x+1)^2} \right\| = \frac{1}{4}.$$

Since this system is very simple it is also possible to calculate the actual maximum which is $\sup_{t\geq 0} \|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| = 1/e$.

As demonstrated above, Theorem 4 is useful if the relative degree of the system transfer function is greater than 1. Now we show a case where where we cannot use this theorem.

Example 2
Consider the same system (21) as before, but now the response to a non-zero initial value $\xi(0) = [x_0, 0]^\top$. *This can be represented by* $\mathcal{B} = [1,0]^T$. *Then we have*

$$\mathcal{C}(sI - \mathcal{A})^{-1}\mathcal{B} = \frac{s+2}{(s+1)^2}.$$

To calculate our first upper bound in (18) we need to encircle the spectrum. The shape here is non-trivial but at least we know that for any circle centered around -1 *with radius smaller than* 1 *we have*

$$\left|\frac{s+2}{(s+1)^2}\right| \leq \left|\frac{2}{(s+1)^2}\right|.$$

The previous calculations give us the upper bound $\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| \leq 2$. *In this case we cannot use Theorem 4 since the relative degree is* 1 *and therefore* $\beta \leq 1$. *However, we can use the very similar Theorem 3 to give an upper bound.* $\|\mathcal{A}\|_\infty = 3$, *and so for any* $R > 3$ *we can use the theorem. It remains to calculate the curve integral along the imaginary axis. Doing this with numerical integration for* $R = 9$ *yields the upper bound* $\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| \lessapprox 2.025$ *(Through optimization this bound can be lowered to* $\|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\| \lessapprox 2.023$*)*

Through these examples we have shown that the best upper bound depends on the situation. The upside of using Theorems 3 and 4 is that they are quite easy to compute numerically. Theorems 1 and 2 relate to the level curves of the input-output pseudospectra and can be qualitatively seen through inspection of these curves, as we will demonstrate in the next section.

4. Application to Networks: Vehicle Strings

To illustrate our bounds, we consider the problem of controlling a string of vehicles—the platooning problem. While performance bounds on platoons and their relation to the network or interaction structure has received ample attention, as we stated in the introduction, the problem calls for bounds relating to the quantity $\sup_t \|y(t)\|_\infty$, where y captures a displacement error. Bounds of this type are important, especially in platooning, since they directly relate to the allowable spacing between consecutive vehicles. However, they tend to be difficult to derive analytically. Here we illustrate how our pseudospectra-inspired approach can be used to evaluate string stability properties for various platoon structures in terms of this quantity.

For this purpose, consider a platoon of size n where each unit is modelled as a double integrator in one spatial dimension, i.e.

$$\ddot{x}_k = u_k$$

where x_k is the position of the kth vehicle with respect to a fix reference and u_k is the input force at vehicle k.

To control the platoon we consider a control law that depends on relative distances to neighboring vehicles and relative to a speed reference. For $k \in \{2,\dots,n-1\}$ we get:

$$\begin{aligned} u_k = (1+\beta_d)(\dot{x}_{k-1}-\dot{x}_k) - (1-\beta_d)(\dot{x}_k-\dot{x}_{k+1}) + \alpha(v_{\text{ref}}-\dot{x}_k) \\ +(1+\beta_p)(x_{k-1}-x_k-d) - (1-\beta_p)(x_k-x_{k+1}-d), \end{aligned} \quad (22)$$

where β_d and β_p are parameters capturing the degree of symmetry in the control law (i.e., look-ahead vs. look-behind control), d a desired intervehicle spacing, v_{ref} is a velocity reference, and $\alpha \geq 0$ is a weight. For the first and last vehicles, we simply define $\ddot{x}_1 = -(1-\beta_d)(\dot{x}_1-\dot{x}_2) + \alpha(v_{\text{ref}}-\dot{x}_1) - (1-\beta_p)(x_1-x_2-d)$, $\ddot{x}_n = (1+\beta_d)(\dot{x}_{n-1}-\dot{x}_n) + \alpha(v_{\text{ref}}-\dot{x}_n) + (1+\beta_p)(x_{n-1}-x_n-d)$. By considering the translated dynamics $\hat{x}_k = x_k + kd$ we get the same dynamics as if we assume $d=0$, so for simplicity we set $d=0$ and consider the dynamics around this equilibrium.

The closed-loop system can be written:

$$\begin{aligned} \begin{bmatrix} \dot{x} \\ \ddot{x} \end{bmatrix} &= \begin{bmatrix} 0 & I \\ -L_p & -L_d - \alpha I \end{bmatrix} \begin{bmatrix} x \\ \dot{x} \end{bmatrix} + \begin{bmatrix} 0 \\ \alpha \mathbf{1} \end{bmatrix} v_{\text{ref}} = \mathcal{A}\xi + \mathcal{B}v_{\text{ref}}. \\ y &= \begin{bmatrix} C & 0 \end{bmatrix} \begin{bmatrix} x \\ \dot{x} \end{bmatrix} = \mathcal{C}\xi, \end{aligned} \quad (23)$$

where L_p, $L_d \in \mathbb{R}^{n\times n}$ are graph Laplacians capturing the vehicle interactions (see further down for definitions). This and similar systems are well studied, see e.g. [Stüdli et al., 2017].

A way to ensure the platoon is well-behaved (e.g. string stable) is to make α large in comparison to L_d and L_p. This, however, essentially transforms the problem to an open-loop system, which is obviously problematic in a real-world setting with disturbances and measurement noise or bias. This motivates the use of a fairly small α, allowing the inter-vehicle adjustments to dominate. In this example, we will use $\alpha = 0.1$.

We use the framework from Section 3 to analyze this system for two cases, one where both Laplacians are asymmetric (a directed string) and one where both are symmetric (bidirectional string). For each system we consider the output

$$y = \begin{bmatrix} x_1 - x_2 \\ x_{\lfloor N/2 \rfloor} - x_{\lfloor N/2 \rfloor + 1} \\ x_{n-1} - x_n \end{bmatrix},$$

which samples three inter-vehicle distances: at the start, middle, and end of the platoon. We will consider the initial condition response, i.e. $\mathcal{B} = I_{2n}$. We expect a string unstable system to perform poorly for at least one of these outputs.

One merit of our proposed method is the possibility to analyze very large systems. In both cases considered here, we therefore model a platoon of $n = 400$ vehicles. We remark that it would of course be possible to simulate the systems for many different inputs and through simulation bound the possible outputs. But as n grows, this quickly becomes very computationally heavy. Using our theorems generates bounds on the worst-case input without any additional effort.

4.1 Directed vehicle string

Consider the control law (22) with $\beta_p = \beta_d = 1$, which renders it fully asymmetric. In this case, we obtain in (23) $L_p = L_d = L_{\text{asym}}$, with

$$L_{\text{asym}} = \begin{bmatrix} 0 & & & \\ -2 & 2 & & \\ & \ddots & \ddots & \\ & & -2 & 2 \end{bmatrix}. \tag{24}$$

In Figure 1 we show the shape of the input-output pseudospectra corresponding to an initial condition response, that is, the level curves of the quantity $\|\mathcal{C}(sI-\mathcal{A})^{-1}I_{2n}\|_\infty$ of (23) with $n = 400$ vehicles. We can see that the level curves extend far into the right half plane with magnitudes of order 10^{30} where $\text{Re}(s)$ is of order 10^{-1}. Through the lower bound in Theorem 1 we can immediately see that there will be a large transient of $\|y(t)\|_\infty$ in at least the orders of 10^{29}. Through a line search, starting at the maximum along the imaginary axis and going into the right half plane we learn that the lower bound in amplification from the worst-case initial conditions to the output (see Section 2.2) is at least

$$\sup_{t\geq 0} \|y(t)\|_\infty \gtrapprox 4.3 \cdot 10^{31},$$

for some ξ_0 such that $\|\xi_0\|_\infty \leq 1$.

Figure 2 displays a Bode plot of the system for various platoon sizes n. That is, we plot the amplitude $\|\mathcal{C}(sI-\mathcal{A})^{-1}I_{2n}\|_\infty$ for $s = i\omega$, $\omega \in (0,\infty)$. Here, we can see the extreme amplification of the frequency response close to the frequency $\omega = 1$. According to our Theorem 3, we can use this frequency response to calculate an upper bound on the transient through integration. By numerical integration we can estimate the upper bound to be

$$\sup_{t\geq 0} |y(t)\|_\infty \lessapprox 1.4 \cdot 10^{33}.$$

From these two bounds we can already conclude that this topology is not suitable for a string of vehicles.

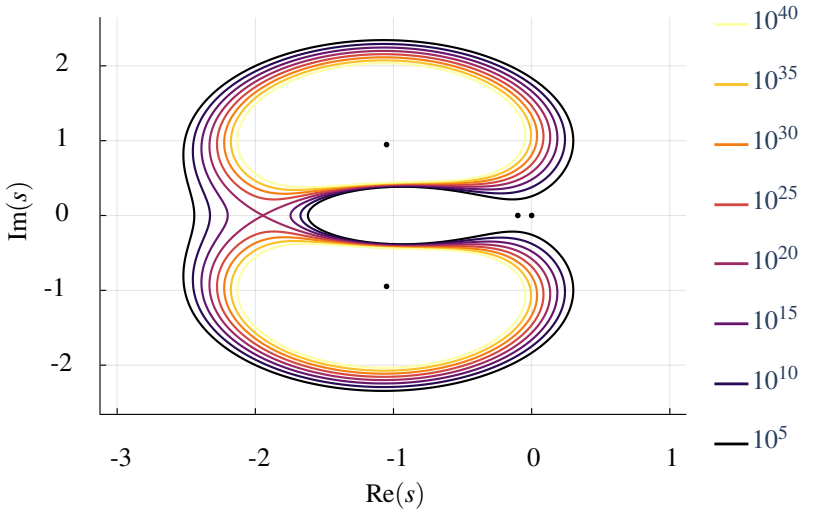

Figure 1. The input-output pseudospectra of (23) for a directed vehicle string with $n = 400$ vehicles. The black dots are the eigenvalues of $\mathcal{A}$. The large values of the input-output pseudospectra even for small s in the right half plane indicate an unfavorable lower bound in Theorem 1.

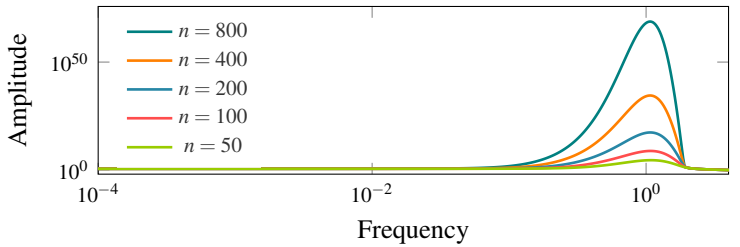

Figure 2. Bode plot displaying the worst-case frequency response of (23) for a directed vehicle string. The amplitude is measured in $\|\cdot\|_\infty$ and the response is shown for various platoon lengths n.

Remark 1
The bounds were possible to compute, since the inversion $\mathcal{C}(sI-\mathcal{A})^{-1}$ can be reduced to the sparse problem $C(s^2I+sL_d+L_p)^{-1}$, where $C \in \mathbb{R}^{3,n}$. As L_d and L_p are tridiagonal this can be computed in $\mathcal{O}(n)$ operations.

4.2 Bidirectional (symmetric) vehicle string

Now, let $\beta_d = \beta_p = 0$ in (22), leading to the symmetric Laplacians $L_p = L_d = L_{\text{sym}}$, with

$$L_{\text{sym}} = \begin{bmatrix} 1 & -1 & & & \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ & & & -1 & 1 \end{bmatrix}. \tag{25}$$

This corresponds to a bidirectional string of vehicles. In Figure 3 we show the shape of the input-output pseudospectra corresponding to an initial condition response of (23) with $n = 400$ vehicles. We can see that the input-output spectra is quite well-behaved and do not extend far into the right half plane, indicating transients will be modest. By Theorem 1 and a line search along the real axis, we learn that the lower bound in amplification from initial conditions to output is

$$\sup_{t\geq 0} \|y(t)\|_\infty \gtrapprox 2.2,$$

for some ξ_0 such that $\|\xi_0\|_\infty \leq 1$.

In Figure 4 we can see the frequency response calculated for various n. Interestingly, the common slope among the curves seems to only behave like a square root which would mean that there is an upper bound independent of the platoon length which bounds the transients due to arbitrary non-zero initial conditions. The numerical upper bound when $n = 400$ was calculated to

$$\sup_{t\geq 0} |\mathcal{C}e^{t\mathcal{A}}\xi_0\|_\infty \lessapprox 9.3.$$

This can be compared with the upper bound calculated for $n = 10^6$ which evaluated at 9.4.

5. Conclusions

In this work we have proposed a pseudospectra-based approach to analyze transient performance of input-output systems, and generalized existing bounds for this purpose. Through our bounds it is possible to quantify $\sup_{t\geq 0} \|\mathcal{C}e^{t\mathcal{A}}\mathcal{B}\|$ in terms of lower and upper bounds. These can be seen as bounding the performance from a worst-case input disturbance to an output $y(t)$ in any p-norm—otherwise often intractable to study. Regarding the problem of controlling vehicle strings in Section 4, we illustrated one application where we believe our bounds can be useful, opening the door to future analysis. For instance, deriving analytical bounds for special network structures. The theorems can also be used to numerically calculate bounds for network structures where the worst inputs are non-obvious, for instance when the agents are non-homogeneous or interaction matrices non-normal.

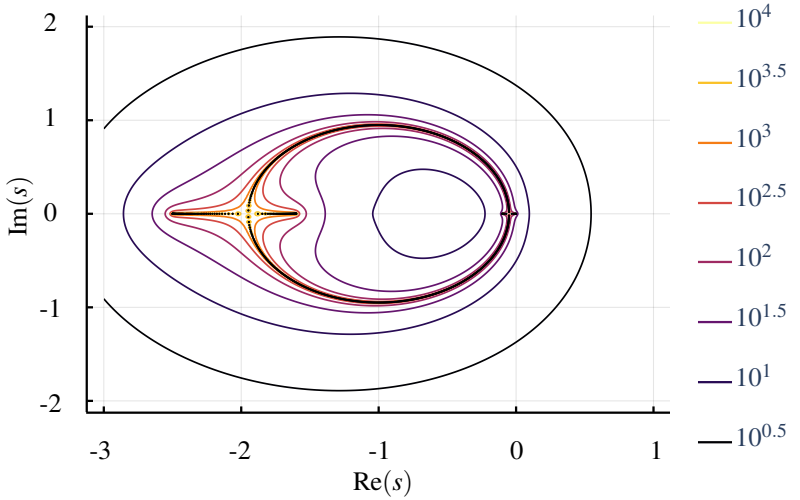

Figure 3. The input-output pseudospectra of (23) for a bidirectional vehicle string with $n = 400$ vehicles. The black dots are the eigenvalues of $\mathcal{A}$. The level curve corresponding to $\|\mathcal{C}(sI - \mathcal{A})^{-1}\| = 10^{1.5}$ can be roughly inscribed in a 1-radius circle, which hints, through Theorem 2, that the transients of $\|y(t)\|_\infty$ will be small.

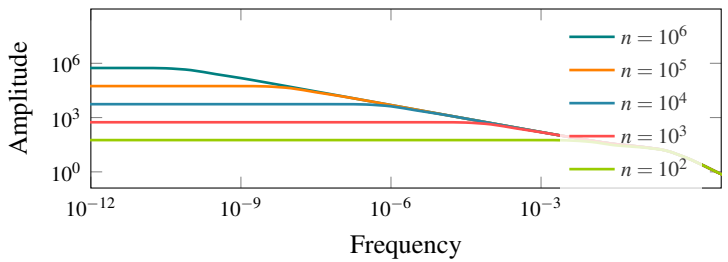

Figure 4. Bode plot displaying the worst-case frequency response of (23) for a bidirectional vehicle string. The amplitude is measured in $\|\cdot\|_\infty$ and the response is shown for various platoon lengths n.

Acknowledgements

The authors are with the Department of Automatic Control and the ELLIIT Strategic Research Area at Lund University, Lund, Sweden. This work was partially funded by Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation and the Swedish Research Council through Grant 2019-00691.

References

Bamieh, B., M. R. Jovanović, P. Mitra, and S. Patterson (2012). "Coherence in large-scale networks: Dimension-dependent limitations of local feedback". *IEEE Trans. Autom. Control* **57**:9, pp. 2235–2249.

Chu, K. (1974). "Decentralized control of high-speed vehicular strings". *Transp. Sci.* **8**:4, pp. 361–384.

Fax, J. A. and R. M. Murray (2004). "Information flow and cooperative control of vehicle formations". *IEEE Trans. Autom. Control* **49**:9, pp. 1465–1476.

Feintuch, A. and B. Francis (2012). "Infinite chains of kinematic points". *Automatica* **48**:5, pp. 901–908.

Feng, S., Y. Zhang, S. E. Li, Z. Cao, H. X. Liu, and L. Li (2019). "String stability for vehicular platoon control: definitions and analysis methods". *Ann. Rev. Control* **47**, pp. 81–97.

Green, K., A. Champneys, and M. Friswell (2006). "Analysis of the transient response of an automatic dynamic balancer for eccentric rotors". *Int. J. Mech. Sci.* **48**:3, pp. 274–293.

Herman, I., S. Knorn, and A. Ahlén (2017). "Disturbance scaling in bidirectional vehicle platoons with different asymmetry in position and velocity coupling". *Automatica* **82**, pp. 13–20.

Kreiss, H.-O. (1962). "Über die stabilitätsdefinition für differenzengleichungen die partielle differentialgleichungen approximieren". *BIT Numer. Math.* **2**, pp. 153–181.

Lancaster, P. and P. Psarrakos (2005). "On the pseudospectra of matrix polynomials". *SIAM J. Matrix Anal. Appl.* **27**:1, pp. 115–129.

Levine, W. and M. Athans (1966). "On the optimal error regulation of a string of moving vehicles". *IEEE Trans. Autom. Control* **11**:3, pp. 355–361.

Matsuo, T. (1994). "Some transfer-function conditions for a desired maximum amplitude or exponential envelope of a closed-loop transient response". In: *IEEE Conf. Decis. Control*. Vol. 3, pp. 2659–2660.

Olfati-Saber, R. and R. M. Murray (2004). "Consensus problems in networks of agents with switching topology and time-delays". *IEEE Trans. Autom. Control* **49**:9, pp. 1520–1533.

Plischke, E. (2005). *Transient Effects of Linear Dynamical Systems*. PhD thesis. University of Bremen.

Ploeg, J., N. van de Wouw, and H. Nijmeijer (2014). "Lp string stability of cascaded systems: application to vehicle platooning". *IEEE Trans. Control Syst. Technol.* **22**:2, pp. 786–793.

Seiler, P., A. Pant, and K. Hedrick (2004). "Disturbance propagation in vehicle strings". *IEEE Trans. Autom. Control* **49**:10, pp. 1835–1842.

Siami, M. and N. Motee (2016). "Fundamental limits and tradeoffs on disturbance propagation in large-scale dynamical networks". *IEEE Trans. Autom. Control* **61**:12, pp. 4055–4062.

Stüdli, S., M. Seron, and R. Middleton (2017). "From vehicular platoons to general networked systems: string stability and related concepts". *Annu. Rev. Control* **44**, pp. 157–172.

Swaroop, D. and J. Hedrick (1996). "String stability of interconnected systems". *IEEE Trans. Autom. Control* **41**:3, pp. 349–357.

Tegling, E., P. Mitra, H. Sandberg, and B. Bamieh (2019). "On fundamental limitations of dynamic feedback control in regular large-scale networks". *IEEE Trans. Autom. Control* **64**:12, pp. 4936–4951.

Tisseur, F. and N. J. Higham (2001). "Structured pseudospectra for polynomial eigenvalue problems, with applications". *SIAM J. Matrix Anal. Appl.* **23**:1, pp. 187–208.

Trefethen, L. N. and M. Embree (2005). *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*. Princeton University Press, Princeton.

Skalbar design för reglering av nätverkssystem: Koordinering med en grannes perspektiv

Jonas Hansson

Institutionen för Reglerteknik

Populärvetenskaplig sammanfattning av doktorsavhandlingen *Scalable Control Design for Networked Systems: Coordination through Local Cooperation*, Juni 2025. Avhandlingen kan laddas ner från: http://www.control.lth.se/publications

---

Denna avhandling undersöker nya metoder för att koordinera stora nätverk av samverkande system, såsom fordon i trafik eller drönare i luften. Målet är att skapa robusta och effektiva formationer, även när systemen växer och blir mer komplexa.

Vi har alla hört talas om smarta bilar; fordon som koordinerar med sina uppkopplade grannar för att smidigt ta sig fram i städer och på motorvägar. Trots att tekniken för sådan kommunikation och koordinerade manövrar existerar, är den ännu inte standard i dagens fordon – varför? En orsak är svårigheten att skapa koordineringsprotokoll som fungerar både för stora och små formationer. Om särskild hänsyn inte tas till hur fordon koordineras kan svängningar uppstå och förstärkas genom långa bilformationer. Dessa svängningar sprider sig likt en ljudvåg eller ett dragspel genom formationen.

En svärm av drönare som koordinerar för att skapa formen av ett hjärta. Bilden är genererad med hjälp av DALL-E.

Effekten är tydlig när långa bilköer bildas vid stoppljus eller efter olyckor. Vid en olycka kan det exempelvis råda fritt flöde precis där olycksplatsen passeras, medan bilar längre bak växelvis startar och stoppar. Det är lätt att tro att detta enbart beror på mänskliga faktorer, vilket bara delvis är sant. Samma problem uppstår med moderna koordineringssystem om de används i tillräckligt långa formationer. Små variationer längst fram i formationen kan då ge upphov till stora svängningar längre bak, vilket till slut tvingar formationen att brytas.

I mitt arbete undersöker jag hur relativa hastigheter och positioner kan användas för att skapa formationer som undviker dessa problematiska svängningar. Lösningen jag föreslår kallas *seriekonsensus*, och idén bakom är relativt enkel. I stället för att direkt använda relativa hastigheter och positioner för att koordinera har jag utvecklat en lösning där två enklare konsensusprotokoll utförs i serie (därav namnet!). I exemplet bilformation kan den nya strategin realiseras genom att varje bil auto-

matiskt, med hjälp av radar och kamera, försöker matcha hastigheten med bilen framför, samtidigt som de försöker hjälpa densamme att positionera sig mitt mellan sina grannar. Detta enkla perspektivskifte förhindrar effektivt små fel från att växa till storskaliga svängningar.

En bilformation som använder seriekonsensus för koordinering. Utöver mätningar av relativa avstånd och hastigheter till närmsta granne tar varje bil även hänsyn till sina grannars uppmätta relativa avstånd.

Teorin bakom seriekonsensus är även relevant bortom koordination av bilformationer. Liknande fenomen finns i många andra naturliga och tekniska system: fågelflockar som flyger i v-formation, fiskstim, drönarsvärmar, koordinerad värmereglering och frekvensreglering i elnät. Gemensamt för alla dessa system är att varje enhet endast kan använda lokala mätsignaler, utan direktinformation om vad övriga enheter gör. Därför behövs även här en skalbar lösning, där varje deltagare kan utforma sin egen styrlag oberoende av formationens storlek och form, samtidigt som systemets totala prestanda måste förbli god.

Implementeringen av seriekonsensus kräver dock något ökad lokal kommunikation, vilket kan ses som en nackdel. Samtidigt visar jag i avhandlingen att denna extra kommunikation möjliggör robustare styrning och eliminerar i princip helt de stora dragspelseffekterna som annars kan uppstå i stora formationer. Ett av avhandlingens huvudresultat är just att detta är möjligt trots att varje enhet enbart använder lokala mätsignaler. Dessutom klarar seriekonsensus olika typer av osäkerheter, såsom tidsfördröjningar och naturliga variationer, som exempelvis att bilar inte accelererar lika snabbt i uppförsbacke som på plan väg.

Sammanfattningsvis introducerar mitt arbete en ny strategi för robust koordinering av stora formationer. Mina resultat visar att effektiv koordinering kan uppnås genom att alla enheter inkluderar sina närmsta medtrafikanters perspektiv i sin koordineringsstrategi. Med andra ord, för bästa resultat räcker det inte att varje deltagare optimerar för sig själv – bäst blir det när alla samarbetar och inkluderar sina grannars perspektiv; en insikt lika enkel som kraftfull.

Printed by Media-Tryck, Lund 2025 NORDIC SWAN ECOLABEL 3041 0903

Department of Automatic Control
P.O. Box 118, 221 00 Lund, Sweden
www.control.lth.se

PhD Thesis TFRT-1149
ISBN 978-91-8104-543-7
ISSN 0280–5316